# THE HYPERBOLIC GEOMETRY OF RANDOM TRANSPOSITIONS ${ }^{1}$ 

BY NATHANAËL BERESTYCKI<br>University of British Columbia and Ecole Normale Supérieure, Paris

Turn the set of permutations of $n$ objects into a graph $G_{n}$ by connecting two permutations that differ by one transposition, and let $\sigma_{t}$ be the simple random walk on this graph. In a previous paper, Berestycki and Durrett [In Discrete Random Walks (2005) 17-26] showed that the limiting behavior of the distance from the identity at time $c n / 2$ has a phase transition at $c=1$. Here we investigate some consequences of this result for the geometry of $G_{n}$. Our first result can be interpreted as a breakdown for the Gromov hyperbolicity of the graph as seen by the random walk, which occurs at a critical radius equal to $n / 4$. Let $T$ be a triangle formed by the origin and two points sampled independently from the hitting distribution on the sphere of radius an for a constant $0<a<1$. Then when $a<1 / 4$, if the geodesics are suitably chosen, with high probability $T$ is $\delta$-thin for some $\delta>0$, whereas it is always $O(n)$-thick when $a>1 / 4$. We also show that the hitting distribution of the sphere of radius an is asymptotically singular with respect to the uniform distribution. Finally, we prove that the critical behavior of this Gromov-like hyperbolicity constant persists if the two endpoints are sampled from the uniform measure on the sphere of radius an. However, in this case, the critical radius is $a=1-\log 2$.

1. Introduction. Let $\delta_{n}$ be the set of permutations of $\{1,2, \ldots, n\}$, and let $\sigma_{t}$ be the continuous-time random walk on $S_{n}$ that results when randomly chosen transpositions are performed at rate 1 . Let $d\left(\sigma_{t}\right)$ be the distance from the identity $I$ at time $t$, that is, the minimum number of transpositions needed to return to $I$. In a previous paper, Berestycki and Durrett [3] showed

THEOREM 0. As $n \rightarrow \infty, d\left(\sigma_{c n / 2}\right) / n \rightarrow_{p} u(c)$ where

$$
\begin{equation*}
u(c)=1-\sum_{k=1}^{\infty} \frac{1}{c} \frac{k^{k-2}}{k!}\left(c e^{-c}\right)^{k} \tag{1.1}
\end{equation*}
$$

Although it is not easy to see from the formula, the function $u(c)=c / 2$ for $c \leq 1$ and is $<c / 2$ for $c>1$.

[^0]We can think of $\sigma_{t}$ as a random walk on the graph $G_{n}$ with vertices $\ell_{n}$ and edges connecting two permutations that differ by one transposition, so that $G_{n}$ is the Cayley graph of $\ell_{n}$ associated with the set of generators $S=$ \{all transpositions $\}$. Theorem 0 was proved by establishing a connection with Erdős-Renyi random graphs. The phase transition observed for $\sigma_{t}$ is then related to the well-known double jump of the size of connected components of $G(n, c / n)$ at $c=1$. [Here and in all that follows, $G(n, p)$ denotes the Erdős-Renyi random graph with parameters $n$ and $p$, i.e., a random graph on $n$ vertices where each edge is present independently of the others with probability $p$.] We refer the reader to Janson et al. [8] for this and other facts about Erdős-Renyi random graphs.

In this paper we try to investigate some of the geometric implications of Theorem 0 . We find a new connection between the speed of a random walk and the Gromov hyperbolicity of the space in which the random walk is evolving.

Organization of the paper. In Sections 1.1, 1.2, 1.3 we present our results. The proofs of these results can be found successively in Sections 2-8. Each proof is preceded by a restatement of the corresponding theorem for convenience, and by an informal proof which outlines the main ideas used.
1.1. Asymptotic hyperbolicity. The notion of hyperbolicity for a discrete structure such as a group is a notion that goes back to Gromov [7]. As there is no derivative, and thus no curvature available in a discrete space, the idea is to define what hyperbolic means using only elementary properties of the space.

One way to do this is as follows. Let $(X,|\cdot|)$ be a metric space, where $|x-y|$ denotes the distance between $x$ and $y$. For points $x, y$ and $p$ in $X$, define the Gromov inner product by

$$
2(x \mid y)_{p}=|x-p|+|y-p|-|x-y| .
$$

$(x \mid y)_{p}$ thus measures how well the union of the geodesic segments $[p, x] \cup[p, y]$ approximates a geodesic between $x$ and $y$. Gromov's original definition of hyperbolic spaces is as follows. Call $X \delta$-hyperbolic if

$$
\begin{equation*}
(x \mid z)_{p} \geq(x \mid y)_{p} \wedge(y \mid z)_{p}-\delta \tag{1.2}
\end{equation*}
$$

for all $x, y, z$ and $p$. This definition is not very intuitive at first, but fortunately there is an equivalent definition, which can be formulated using the notion of $\delta$-thin triangle. A triangle ( $x, y, z$ ) with geodesic sides $s_{1}, s_{2}, s_{3}$ is said to be $\delta$-thin if any side, say $s_{1}$, lies entirely within distance at most $\delta$ of the two remaining sides:

$$
s_{1} \subset\left\{x \in X, d\left(x, s_{2} \cup s_{3}\right) \leq \delta\right\} .
$$

The space is called $\delta$-hyperbolic if all geodesic triangles are $\delta$-thin, and it is simply called hyperbolic if it is $\delta$-hyperbolic for some $\delta \geq 0$ (when $\delta=0$, the space isometrically embeds into a tree). It is not immediate, but not hard to check, that if
all triangles $(x, y, z)$ are $\delta$-thin, then (1.2) is satisfied for some number $\delta^{\prime}$ that may differ by a constant factor from $\delta$. Conversely, in a space where (1.2) is satisfied for all points ( $p, x, y, z$ ), all triangles are $\delta^{\prime}$-thin, where $\delta^{\prime}$ may differ from $\delta$ by a constant factor.

Of course a bounded space (in particular, a finite space such as $\delta_{n}$ ) is trivially hyperbolic, but we will be interested in situations where the constant $\delta$ may or may not stay bounded as the size of the space tends to $\infty$.

Our first result makes the connection between Theorem 0 and Gromov hyperbolic spaces, where we look at the two definitions of hyperbolic constants suitably weakened. For $0<a<1$, let $\partial B(a n)$ be the sphere of radius $a n$, that is, the set of points at distance $\lfloor a n\rfloor$ from the origin. We let $v$ be the hitting distribution of $\partial B(a n)$ by $\sigma_{t}$, that is, $v$ is the law on $\partial B(a n)$ of $\sigma_{T}$ where $T=\inf \left\{t>0, d\left(\sigma_{t}\right)=\lfloor a n\rfloor\right\}$.

THEOREM 1. Let $x, y$ be sampled from $v$ independently, and set $p=I$, the identity element.

1. If $a<1 / 4$, then there is some $\delta<\infty$ (depending only on $a$ ), such that

$$
E(x \mid y)_{p} \leq \delta
$$

Moreover, with probability asymptotically 1, there is a geodesic between $x$ and $y$ that comes within expected distance $\delta^{\prime}<\infty$ of $p$.
2. If $a>1 / 4$, then

$$
E(x \mid y)_{p} \sim \delta n
$$

for some $0<\delta<\infty$. Moreover, no geodesic between $x$ and $y$ can approach $p$ closer than $\delta^{\prime} n$ with probability asymptotically 1 , where $0<\delta^{\prime}<\infty$.

In the statement of the theorem and in the rest of the paper, $a_{n} \sim b_{n}$ means that $a_{n} / b_{n} \rightarrow 1$.

REMARK. It follows immediately from Theorem 1 that when $a<1 / 4$, with probability asymptotically 1 ,

$$
(x \mid z)_{p} \geq(x \mid y)_{p} \wedge(y \mid z)_{p}-\delta
$$

for independent $x, y, z$ sampled from $v$, hence the idea that definition 1 of hyperbolicity is satisfied "asymptotically $v$-almost surely." The statement about the geodesics shows that definition 2 is satisfied "asymptotically $v$-almost surely" when $a<1 / 4$.

At this point we should emphasize that the result in Theorem 1 involves hyperbolic constants that are different from the standard definitions discussed above in several important ways. The most obvious difference comes from the randomness
of $x$ and $y$, and from the fact that the roles played by $x, y$ and $p$ are somewhat different. Here $p$ is a fixed reference point, whereas Gromov's definition requires that every triangle should be thin. Another issue is that, corresponding to the second definition of hyperbolicity with thin triangles, we show that there exists a certain geodesic between $x$ and $y$ having the desired properties. As we will see below in Theorem 6, there may be a great many geodesics between two given points in $\left\{_{n}\right.$. More importantly, these geodesics can be far apart, as will show the following concrete example:


Since for any permutation $\pi$ we have $d(\pi)=n-\#$ cycles of $\pi, d(\sigma)=8$. $\pi_{1}$ and $\pi_{2}$ are on two geodesics from $I$ to $\sigma$, but $d\left(\pi_{1}, \pi_{2}\right)=d\left(\pi_{1} \pi_{2}^{-1}\right)=8$. In general if $d(\sigma)=c n / 2$ with $c<1$, and we divide the cycles at random into two groups, we can define $\pi_{1}$ to have cycle structure given by the first group of $\sigma$ staying as it is and the second completely broken in cycles on lengths 1 . If we define $\pi_{2}$ by the exchange of the two groups, then we will have $d\left(\sigma, \pi_{i}\right)=c n / 4$ and $d\left(\pi_{1}, \pi_{2}\right)=c n / 2$.
1.2. The geometry of $G_{n}$. How much can we learn from Theorem 1 about the global geometry of $G_{n}$ ? To answer this question, we need to see how special a choice it is to sample the points $x$ and $y$ according to the hitting distribution $v$. (The fact that $p=I$ is a fixed reference point is not too important, due to the transitivity of $G_{n}$.) We begin by an apparently unrelated question, which is to ask how large is a ball of radius $a n$.

THEOREM 2. If $0 \leq a \leq 1$, then as $n \rightarrow \infty$, we have $|B(I, a n)| \approx(n!)^{a}$ in $a$ logarithmic sense, that is,

$$
\lim _{n \rightarrow \infty} \frac{\log |B(I, a n)|}{n \log n}=a
$$

This result is probably not new, but we have not found it in the literature. Our original motivation for studying the volume growth in $G_{n}$ was to try to understand the phase transition of Theorem 0 in terms of the geometry of $G_{n}$. Our first thought was that since the speed was nonsmooth we might see a change in the volume growth. The above result contradicts this idea.

To put our next two results into perspective it is useful to contrast them with Brownian motion $B_{t}$ on a $d$-dimensional manifold of constant negative curvature -1 . In that case as $t \rightarrow \infty$, if $d\left(B_{t}\right)$ is the distance from the origin, then (see [11], e.g.) there is a constant $v$ so that

$$
d\left(B_{t}\right) / t \rightarrow v \quad \text { as } t \rightarrow \infty
$$

In the case of Brownian motion on hyperbolic space, rotational symmetry implies that the hitting distribution is uniform. In contrast for the random transposition random walk, we will see in Theorem 3 that the hitting distribution is asymptotically singular with respect to the uniform distribution on $\partial B(I, a n)$.

THEOREM 3. Let $\left|\mathcal{C}_{1}\right|$ be the length of the cycle that contains 1 . Under $\mu$, the uniform distribution on $\partial B(I$, an $)$,

$$
\left|\mathcal{C}_{1}\right| \Rightarrow \mathbf{G}
$$

where $\mathbf{G}$ is a geometric r.v. with $P(\mathbf{G}>k)=(b /(1+b))^{k}$ and $b$ satisfies $\log (1+$ b) $/ b=1-a$.

To describe the hitting distribution $v$, we note that (1.1) suggests that it will be the same as the distribution of $\sigma_{c n / 2}$ where $c=u^{-1}(a)$. When $a>1 / 2$ this is much different from the distribution in Theorem 3 since in this case $c>1$ and Schramm [12] has shown that $\sigma_{c n / 2}$ has cycles of lengths of order $n$.

Here we will concentrate on what happens when $a<1 / 2$ and $c=2 a$. In this case results in [3] show that as $n \rightarrow \infty$, the number of fragmentations before time $c n / 2$ is asymptotically a Poisson random variable with mean $\kappa(c)=$ $-(\log (1-c)+c) / 2$. In particular,

$$
P\left(d\left(\sigma_{c n / 2}\right)=c n / 2\right) \rightarrow e^{-\kappa(c)}=e^{c / 2} \sqrt{1-c}
$$

It will be convenient to approach the hitting distribution $\nu$ by the distribution $\nu_{0}$ of $\sigma_{c n / 2}$ conditioned on no fragmentation. More generally, if $v_{k}=v$ conditioned on exactly $k$ fragmentations before the hitting time,

$$
\nu=e^{-\kappa(c)} \sum_{k=0}^{\infty} v_{k} \frac{\kappa(c)^{k}}{k!}+o(1)
$$

To study $\nu_{0}$, we recall the connection with random graphs developed in [3]: when we transpose $i$ and $j$ we draw an edge between $i$ and $j$. In order for the distance from the identity to increase by 1 at each time, each transposition must involve indices from two different cycles and will merge them into one. In terms of the random graph, this means that all components are trees. Using results from [3], it is straightforward to show:

THEOREM 4. Let $\mathcal{C}_{1}$ be the length of the cycle that contains 1 . Let $c<1$. Under $\nu_{0}$,

$$
P\left(\left|\mathbb{C}_{1}\right|=k\right) \rightarrow \frac{1}{c} \frac{k^{k-1}}{k!}\left(c e^{-c}\right)^{k} \quad \text { for all } k \geq 1
$$

Theorems 3 and 4 show that the uniform distribution $\mu$ and the hitting distribution $\nu_{0}$ concentrate on different permutations. In the first case the number of fixed points will be close to its expected value $n P\left(\left|\mathcal{C}_{1}\right|=1\right)=n /(1+b)$. In the second it will be close to $n e^{-c}$ by Theorem 4. This is made precise by the following theorem.

THEOREM 5. As $n \rightarrow \infty$, the hitting distribution $v$ and the uniform distribution $\mu$ on a sphere of radius an are asymptotically singular:

$$
d_{\mathrm{TV}}(\mu, v) \rightarrow 1
$$

Let $t=[c n / 2]$ with $c<1$. To understand why $v$ is different from $\mu$ we will examine the Radon-Nikodym derivative $r(\sigma)=d \nu_{0} / d \mu$. It is not hard to show that

THEOREM 6. Suppose $d(\sigma)=t$ and $m_{1}, \ldots, m_{j}$ are the cycle lengths of $\sigma$. The number of paths of length $t$ from $I$ to $\sigma$ is

$$
t!\prod_{i=1}^{j} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!} .
$$

If $t=\lfloor c n / 2\rfloor$ with $c<1$, then

$$
r(\sigma)=K_{n, t} \prod_{i=1}^{j} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!},
$$

where $K_{n, t}$ is a constant that only depends on $n$ and $t$.
The last result enables us to prove a stronger version of Theorem 5: it tells us that the "support" of $v$ is concentrated on a set that is exponentially smaller than the size of $\partial B(a n)$.

THEOREM 7. Suppose $a<1 / 2$. There exists a set $S_{n} \in \partial B(a n)$ such that $\nu\left(S_{n}\right) \rightarrow 1$ as $n \rightarrow \infty$ and

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \log \frac{\left|S_{n}\right|}{\mid \partial B(\text { an }) \mid}=\gamma<0 .
$$

1.3. The hyperbolic constant under the uniform measure. In Theorem 1, we learn that if $x$ and $y$ are sampled from $v$, roughly speaking, the Gromov hyperbolicity of the "support" breaks down at $a=1 / 4$, that is, the hyperbolic constant increases suddenly from $O(1)$ to $O(n)$ at this point.

However, the results from the previous section tell us that this "support" is (exponentially) small with respect to the ambient space. It is therefore natural to ask what happens to Theorem 1 when we replace $\nu$ with the uniform measure $\mu$ on
$\partial B(a n)$. Theorem 8 will show that the qualitative behavior of the hyperbolic constant remains the same. We prove that there is a threshold where the expected Gromov inner product $E(\sigma \mid \pi)_{p}$ jumps from $O(1)$ to $O(n)$, but this time the critical value is $a=1-\log 2 \approx 0.31$, rather than $a=1 / 4$.

When $\sigma$ and $\pi$ are independent uniform permutations on $\partial B(a n)$, by the transitivity of $G_{n}$, it is enough to analyze $d(\sigma, \pi)$ to understand $(\sigma \mid \pi)_{p}$, the inner Gromov product. Since $d(\sigma, \pi)=d\left(I, \sigma^{-1} \pi\right)$, which has the same law as $d(I, \sigma \pi)$, it will be enough to characterize the values of $a$ for which $d(I, \sigma \pi)=$ $2 a n+o(n)$ and those for which it is $<2 a n$.

THEOREM 8. Let $0<a<1$ and let $\sigma, \pi$ be two random independent points chosen uniformly from $\partial B(a n)$. Then:

1. If $a<1-\log 2$,

$$
E(\sigma \mid \pi)_{p} \leq \delta(\log n)^{2}
$$

for some $0<\delta=\delta(a)<\infty$. Moreover, with probability asymptotically 1 , there is a geodesic between $\sigma$ and $\pi$ that comes within distance at most $\delta(\log n)^{2}$ of $p$.
2. If $a>1-\log 2$,

$$
E(\sigma \mid \pi)_{p} \sim \delta n
$$

for some $\delta=\delta(a)>0$. Moreover, no geodesic can approach $p$ closer than $\delta^{\prime} n$ for some $0<\delta^{\prime}<\infty$.

REMARK. The $O\left((\log n)^{2}\right)$ bound in part 1 of the theorem could probably be improved into an $O(1)$ bound (just like in Theorem 1) with some more work, but we have not tried to do so. In part 2, by analogy with Berestycki and Durrett [3], we conjecture that the fluctuations are of order exactly $n^{1 / 2}$ in the supercritical regime. More precisely, it should be true that when $a>1-\log 2$,

$$
n^{-1 / 2}\left(E(\sigma \mid \pi)_{p}-\delta n\right) \Rightarrow \mathcal{N}(0, \kappa)
$$

where $\delta$ is the limit in part 2 of the theorem, and $\kappa$ is some parameter.
2. Asymptotic hyperbolicity under v. The first result we prove is Theorem 1 .

THEOREM 1. Let $x, y$ be sampled from $v$ independently, and set $p=I$, the identity element.

1. If $a<1 / 4$, then there is some $\delta<\infty$ (depending only on $a$ ), such that

$$
E(x \mid y)_{p} \leq \delta
$$

Moreover, with probability asymptotically 1, there is a geodesic between $x$ and $y$ that comes within expected distance $\delta^{\prime}<\infty$ of $p$.
2. If $a>1 / 4$, then

$$
E(x \mid y)_{p} \sim \delta n
$$

for some $0<\delta<\infty$. Moreover, no geodesic between $x$ and $y$ can approach $p$ closer than $\delta^{\prime} n$ with probability asymptotically 1 , where $0<\delta^{\prime}<\infty$.

Sketch of the proof. Let $X_{t}$ and $Y_{t}$ be two independent random walks starting at the origin. Let them run until the times $T$ and $T^{\prime}$ where they respectively hit the sphere $\partial B(a n)$. Then the transitivity of the Cayley graph of $\delta_{n}$, and the reversibility of the increments of the random walk, imply that $\left(X_{T}, X_{T-1}, \ldots, p, Y_{1}, \ldots, Y_{T^{\prime}}\right)$ is a random walk path of length $T+T^{\prime}$. Hence the distance between $X_{T}=x$ and $Y_{T^{\prime}}=y$ is the same as $d\left(\sigma_{T+T^{\prime}}\right)$. By Theorem $0, T$ and $T^{\prime} \approx \frac{1}{2} u^{-1}(a) n$, so applying Theorem 0 again, when $a<1 / 4$, $|x-y| \approx 2 a n=|x|+|y|$ [the random walk runs for a time $2 a n<n / 2$ and there are only $O(1)$ fragmentations]. For $a>1 / 4$, the random walk is run for time $u^{-1}(a) n$ which, in view of Theorem 0 , means that $c=2 u^{-1}(a)$, and $|x-y|=n u\left(2 u^{-1}(a)\right) \ll 2 a n$. See Figure 1.

The claim about the existence of a geodesic that makes the triangle $(x, p, y)$ thin involves necessarily another argument, since geodesics may be far apart. However, it is not very hard to construct by hand a geodesic between the identity and $x$ such that each point of the random walk path is within $O(1)$ of this geodesic. Applying this construction to the two random walk paths gives the result of Theorem 1.

Proof of Theorem 1. Let us first deal with case $a<1 / 4$ and prove that in this case $E(x \mid y)_{p} \leq \delta$. Keeping the same notation as above, note that $a n+O$ (1) steps are sufficient for $X$ to reach distance $a n$. Indeed, after an steps, $X_{a n}$ is at distance $a n-X_{1}$ where $X_{1}$ is twice a Poisson random variable by Theorem 1 in [3]. It is immediate that in $X_{1}$ steps the probability that $X$ has a fragmentation converges to 0 . Therefore $T-a n \Rightarrow X_{1}$ (remember that here time


FIG. 1. Two independent random walks run until they hit the sphere of radius an.
is measured discretely). Similarly $T^{\prime}-a n \Rightarrow X_{2}$ where $X_{1}, X_{2}$ are i.i.d. Hence $\left(x=X_{t}, X_{T-1}, \ldots, X_{1}, I, Y_{1}, \ldots, Y_{t}=y\right)$ is a random walk of $2 a n+X_{1}+X_{2}$ steps. In the worst case possible all $X_{1}+X_{2}$ steps represent "backward" steps (meaning, toward $x$ rather than $y$ ). Hence if $X_{3}=a n-d\left(I, \sigma_{2 a n}\right)$ (so that $X_{3}$ is also twice a Poisson random variable, but with a different parameter),

$$
\begin{aligned}
2 E(x \mid y)_{p} & =2 a n-|x-y| \\
& \leq 2 a n-E d\left(I, \sigma_{2 a n}\right)+E\left(X_{1}+X_{2}\right) \\
& \leq E\left(X_{1}\right)+E\left(X_{2}\right)+E\left(X_{3}\right)<\infty .
\end{aligned}
$$

It is slightly simpler to prove that when $a>1 / 4, E(x \mid y)_{p} \sim \delta n$. Indeed, in this case, by Theorem 0 , we have that

$$
\frac{1}{2} u^{-1}(a)-\varepsilon \leq T / n \leq \frac{1}{2} u^{-1}(a)+\varepsilon
$$

Therefore

$$
\inf _{\left|t / n-u^{-1}(a)\right| \leq 2 \varepsilon} d\left(I, \sigma_{t}\right) \leq|x-y| \leq \sup _{\left|t / n-u^{-1}(a)\right| \leq 2 \varepsilon} d\left(I, \sigma_{t}\right) .
$$

An easy estimate shows that we are never off by more than $O\left(n^{1 / 2}\right)$ if we evaluate the distance of the random walk by counting the number of clusters of the random graph rather than the number of cycles of $\sigma_{t}$. But for the random graph, the number of clusters is monotone increasing. Hence, if $\alpha$ denotes $u\left(2 u^{-1}(a)\right)$, we have by continuity of $u$ that

$$
\alpha-\varepsilon^{\prime}+o(1) \leq \frac{E|x-y|}{n} \leq \alpha+\varepsilon^{\prime}+o(1)
$$

and $\varepsilon^{\prime}$ can be made as small as desired by continuity of $u$. Therefore

$$
\frac{E|x-y|}{n} \rightarrow \alpha
$$

It suffices now to prove that $\alpha<2 a$, that is, $u\left(2 u^{-1}(a)\right)<2 a$ or, after change of variable $c=u^{-1}(a)$, it suffices to prove $u(2 c)<2 u(c)$ for all $c>1 / 2$. This fact is a consequence of the sublinearity of $u$ : it will be proved later that $u$ is strictly concave on $[1, \infty)$, from which it follows that $u(c)>u(2 c) / 2$.

We now turn to the part of the theorem that concerns geodesics, and prove that for a random walk ( $X_{t}, t \leq c n / 2$ ) of time-duration $c n / 2$ with $c<1$, there is a geodesic between $\sigma=X_{c n / 2}$ and $I$, that we call $\gamma$, such that

$$
\begin{equation*}
E \sup _{t \leq c n / 2} d\left(X_{t}, \gamma\right)=O(1) \tag{2.1}
\end{equation*}
$$

This shows that when $c<1$ there is a geodesic that stays close to the random walk path. When $a<1 / 4, p=I$ is on the random walk path that leads from $x$ to $y$, so this shows that $E(d(I, \gamma))=O(1)$, as claimed in the theorem. The case $a>1 / 4$ is trivial by the triangle inequality.

Let $\tau_{1}, \ldots, \tau_{N}$ be the sequence of transpositions that are the increments of the random walk path leading to $\sigma$, so that $\sigma=\tau_{1} \ldots \tau_{N}$. Let $\gamma$ be the geodesic between $\sigma$ and $I$ defined by $\gamma_{0}=\sigma, \gamma_{1}=\sigma \tau_{N}, \gamma_{2}=\gamma_{1} \tau_{N-1}, \ldots$, until the first time $t$ such that multiplying $\gamma_{t}$ by $\tau_{N-t}$ would result in a coagulation of two cycles of $\gamma_{t}$. We do not allow this possibility (otherwise $\gamma$ would not be a geodesic), and simply skip $\tau_{N-t}: \gamma_{t+1}=\gamma_{t} \tau_{N-t-1}$. We will see in a moment that this path never backtracks and that it ends at a bounded distance from $I$, to which it will be necessary to add a (bounded) number of steps so that it actually ends at $I$.

Let $n(t)$ be the index of the transposition to be performed at time $t$ on $\gamma_{t}$. Note that we can always write

$$
\gamma_{t}=\tau_{1} \tau_{2} \ldots \tau_{n(t)} \prod_{i \in K_{t}} \tau_{i}
$$

where $K_{t}$ is a set whose size we will show is bounded. Indeed, even when we skip $\tau_{n(t)}$ in $\gamma_{t}$ [so that $n(t) \in K_{t+1}$ ], the following transpositions $\tau_{n(t)-1}, \ldots$ commute with the members of $K_{t}$ with high probability and they can "jump above" the terms in $K_{t}$ and cancel the rest of the transpositions ( $\tau_{1} \ldots \tau_{n(t)-1}$ ).

Lemma 1. For all $t, E\left(\left|K_{t}\right|\right) \leq O(1)$, where $O(1)$ is a constant that depends only on $c<1$. As a consequence, the path ends at bounded distance from the identity and the distance $E \sup _{t \leq c n / 2} d\left(X_{t}, \gamma\right)=O(1)$.

Proof. There are two ways to add a member to $K_{t-1}$ at time $t$. The first one is that performing $\tau_{n(t)}$ will result in a coagulation, so that it is skipped by $\gamma$. The other way is if $\tau_{n(t)}$ does not commute with one of the members of $K_{t-1}$, it stays stuck somewhere in $K_{t}$.

If $\tau_{n(t)}=(i, j)$, we claim that in order for $i$ and $j$ to be in the same cycle of $\gamma_{t}$, $i$ and $j$ must belong to a component of the Erdős-Renyi graph associated with the random walk that contains a cycle at time $c n / 2$. We will prove this in a moment, but if we admit this, then it follows that all transpositions in $K_{t}$ act on vertices that belong to $U(c n / 2)$, the unicyclic components of the random graph at time $\mathrm{cn} / 2$ : if $i \in K_{t}$, then either $\tau_{i}=(i, j)$ yields a coagulation in $\gamma_{t}$, or it does not commute with some member of $(k, l)$ of $K_{t-1}$, in which case $(i, j)$ overlaps with $(k, l)$. By induction, $k, l \in U(c n / 2)$, therefore so are $i$ and $j$.

Let us prove our claim that if $(i, j)$ would yield a coagulation in $\gamma_{t}$, then $i, j \in U(c n / 2)$. Let us observe first that $i$ and $j$ must already be in the same component of the random graph: because $\tau_{n(t)}$ was performed on the random walk, $i$ and $j$ were connected at that point in the random graph and they remain so. If $i$ and $j$ are in different cycles of $\sigma$, then there must have been some ulterior fragmentation in their cycles, so the claim holds. When they are in the same cycle of $\sigma$, then there must be some transposition $\tau_{i}$ with $i \in K_{t}$ such that $i$ and $j$ are in different cycles of $\gamma$ after $\tau_{i}$. Call those cycles $C_{1}$ and $C_{2} . \tau_{i}$ involves two members
$k$ and $l$ of $C_{1} \cup C_{2}$. Moreover the cycle structure of $\gamma$ before $(k, l)$ is performed must be of the form

$$
(k, \ldots, i, \ldots, j, \ldots, l, \ldots)
$$

otherwise $(k, l)$ cannot separate $i$ and $j$ at the next step. Unless $i$ and $j$ belong to a complex component, this implies that the cycle structure of $\sigma$ has the same form. However, this can only happen if $k$ and $l$ were connected to the component of $i$ and $j$ at different times; otherwise the cycle structure would be of the form $(i, \ldots j, \ldots, k, \ldots, l)$ or $(i, \ldots k, \ldots, l, \ldots, j)$. This implies in turn the existence of a cycle in the random graph component of $i$ and $j$ at time $c n / 2$.

From there it follows in a straightforward way that $\left|K_{t}\right| \leq|U(c n / 2)|$ (in a unicyclic component there are as many edges as vertices). It is now standard in the theory of random graphs to show that $|U(c n / 2)|$ is bounded:

$$
\begin{aligned}
E|U(c n / 2)| & \sim\left(\frac{\pi}{8}\right)^{1 / 2} \sum_{k=2}^{\infty}\binom{n}{k} \frac{k^{k-1 / 2}}{k!}\left(\frac{c}{n}\right)^{k}\left(1-\frac{c}{n}\right)^{k(n-k)+\binom{k}{2}-k} k \\
& \sim\left(\frac{\pi}{8}\right)^{1 / 2} \sum_{k=2}^{\infty} \frac{k^{k+1 / 2}}{k!}\left(c e^{-c}\right)^{k}<\infty
\end{aligned}
$$

which completes the proof of the lemma.
Now let $X_{t}$ be a point on the random walk path. Since $\gamma$ tries to perform all $\tau_{i}$ (at reverse), there is a time $s$ such that $n(s)=t$, that is, the next transposition to be examined by $\gamma_{s}$ is $\tau_{t}$. At this time,

$$
\gamma_{s}=\tau_{1} \cdots \tau_{t} \prod_{i \in K_{s}} \tau_{i}
$$

so that $\left|K_{s}\right|$ steps are enough to reach $\gamma_{s}$ from $X_{t}$. Since $E\left(\left|K_{s}\right|\right)<O(1)$ by the lemma, we have proved that

$$
E \sup _{t \leq c n / 2} d\left(X_{t}, \gamma\right) \leq O(1)
$$

and Theorem 1 is proved.
3. Large deviations and volume growth. The goal of this section is to prove Theorem 2, which we restate here for convenience.

THEOREM 2. If $0 \leq a \leq 1$, then as $n \rightarrow \infty$, we have $\mid B(I$, an $) \mid \approx(n!)^{a}$ in a logarithmic sense, that is,

$$
\lim _{n \rightarrow \infty} \frac{\log |B(I, a n)|}{n \log n}=a
$$

Sketch of the proof. The proof of the result is more interesting than the limit. We begin by recalling the dynamics of the Chinese restaurant process (see, e.g., [9]). Customer 1 enters and sits at table 1. At step $i$, customer $i$ enters and starts a new table with probability $1 / i$ or sits to the left of customer $k$ where $k$ is chosen uniformly at random in $\{1, \ldots, i\}$. From the tables we define a permutation $\sigma$ by $\sigma(i)=i$ if customer $i$ is sitting by himself at his table and $\sigma(i)=k$ if $k$ sits to the right of $i$. It is easy to see that this defines a uniform random permutation on $\ell_{n}$, and that the cycle structure is given by listing the individuals at the tables in clockwise order. It is well known that if $\sigma \in 夕_{n}$, then $d(\sigma)=n$ - the number of cycles of $\sigma$. In the Chinese restaurant process construction, let $\zeta_{i}$ be the random variables taking the value 1 if customer $i$ sits at an existing table (and 0 otherwise). The $\zeta_{i}$ 's are independent Bernoulli $(1-1 / i)$ random variables. Recall that if $\sigma$ is a permutation, then $d(\sigma, I)=d(\sigma)$ is $n-\#$ cycles of $\sigma$. Hence, if $\sigma$ is uniformly distributed over $\wp_{n}$, then $d(\sigma)$ has the same distribution as $S_{n}=\sum_{i=1}^{n} \zeta_{i}$.

The $i$ 's where a new cycle starts (i.e., $\zeta_{i}=0$ ) are distributed with the same law as that of the occurrences of records for i.i.d. variables with continuous distribution function (cf. [5], Example 6.2 of Chapter 1). From calculations in that example it follows that $\left(n-S_{n}\right) / \log n \rightarrow 1$ in probability.

Returning to our calculation of the volume of the ball,

$$
|B(I, a n)|=n!P\left(S_{n} \leq a n\right)
$$

for all $0<a<1$. It is straightforward to generalize large deviation results for i.i.d. random variables (see, e.g., [5], Section 2.9) to prove Theorem 1. One begins with the observation that for $\lambda>0$

$$
\begin{equation*}
P\left(S_{n} \leq a n\right) \leq e^{\lambda a n} E e^{-\lambda S_{n}} \tag{3.1}
\end{equation*}
$$

optimizes the upper bound over $\lambda$ and uses a change of measure argument to prove a corresponding lower bound.

Proof of Theorem 2. Let $\left\{\zeta_{i}, i \geq 1\right\}$ be independent with $P\left(\zeta_{i}=1\right)=$ $1-1 / i$, and let $S_{n}=\sum_{i=1}^{n} \zeta_{i}$. Since $(\log n!) /(n \log n) \rightarrow 1$ it suffices to show:

Lemma 2. Let $0<a<1$. As $n \rightarrow \infty$,

$$
\lim _{n \rightarrow \infty} \frac{\log P\left[S_{n} \leq a n\right]}{n \log n}=a-1
$$

Proof. Let $\varphi_{n}(\lambda)=\mathbf{E}\left[e^{-\lambda S_{n}}\right]$. Using the definition we have

$$
\varphi_{n}(\lambda)=\prod_{i=1}^{n}\left[\left(1-\frac{1}{i}\right) e^{-\lambda}+\frac{1}{i}\right] \equiv \prod_{i=1}^{n} q_{i}
$$

where $\equiv$ indicates that the last equation is the definition of $q_{i}$. By Markov's inequality we have

$$
\begin{equation*}
\log P\left[S_{n} \leq a n\right] \leq n\left(\lambda a+\frac{1}{n} \log \varphi_{n}(\lambda)\right) \quad \text { for all } \lambda \tag{3.2}
\end{equation*}
$$

If we define

$$
F_{\lambda}(x)=\frac{1}{\varphi_{n}(\lambda)} \int_{-\infty}^{x} e^{-\lambda y} d F^{n}(y)
$$

then $F_{\lambda}$ is a distribution function such that

$$
\operatorname{mean}\left(F_{\lambda}\right)=-\frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)} \quad \text { and } \quad \operatorname{var}\left(F_{\lambda}\right)=\frac{d}{d \lambda} \frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)} \geq 0
$$

To optimize (3.2), we want to choose $\lambda$ so that

$$
a+\frac{1}{n} \frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)}=0
$$

This says that the mean of the transformed distributions is $n a$, so

$$
a=\frac{1}{n} \sum_{i=1}^{n} \frac{(1-1 / i) e^{-\lambda}}{q_{i}}=\frac{1}{n} \sum_{i=1}^{n} \frac{(i-1) e^{-\lambda}}{(i-1) e^{-\lambda}+1} .
$$

We guess that the optimal $\lambda$ must be given (asymptotically) by $e^{-\lambda_{\text {opt }}}=b / n$. Plugging this in the above gives

$$
a=\frac{1}{n} \sum_{j=1}^{n-1} \frac{j b / n}{(j b / n)+1} \rightarrow \int_{0}^{1} \frac{b x}{b x+1} d x=1-\frac{1}{b} \log (b+1) .
$$

From this we see that we should choose $b$ so that $\log (b+1) / b=1-a$.
Upper bound. Let us calculate what (3.2) gives with this choice of $\lambda$ :

$$
\begin{align*}
\frac{1}{n} \log \varphi_{n}(\lambda) & =-\log n+\frac{1}{n} \log \prod_{i=1}^{n}\left(\left(1-\frac{1}{i}\right) b+\frac{n}{i}\right)  \tag{3.3}\\
& \rightarrow-\log n+\int_{0}^{1} \log (b+1 / x) d x \tag{3.4}
\end{align*}
$$

Since the last integral is finite it follows from (3.2) and $\lambda_{\text {opt }}=-\log b+\log n$ that

$$
\limsup _{n \rightarrow \infty} \frac{1}{n \log n} P\left(S_{n} \leq n a\right) \leq a-1
$$

proving the upper bound half of Lemma 2.

Lower bound. The argument is similar to that in [5], page 73. Fix any $v<a$ and $v<v^{\prime}<a$. Define a real number $b^{\prime}$ by

$$
\frac{\log \left(1+b^{\prime}\right)}{b^{\prime}}=1-v^{\prime}
$$

For any $\lambda$,

$$
\begin{aligned}
P\left[S_{n} \leq a n\right] & \geq \int_{v n}^{a n} d F^{n}(x) \geq \int_{v n}^{a n} e^{\lambda x} \varphi_{n}(\lambda) d F_{\lambda}(x) \\
& \geq \varphi_{n}(\lambda) e^{\lambda n v}\left[F_{\lambda}(n a)-F_{\lambda}(n v)\right]
\end{aligned}
$$

First, we prove that we can choose $\lambda$ such that $\left[F_{\lambda}(n a)-F_{\lambda}(n \nu)\right] \rightarrow 1$. Recall that the mean of $F_{\lambda}$ is $-\frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)}$, and that the latter function starts at $n-\log n$ for $\lambda=0$, is strictly decreasing and equals $n a$ when $\lambda=\lambda_{\mathrm{opt}}=-\log b+\log n$, that is, $e^{-\lambda_{\text {opt }}}=b / n$. Thus if we pick $\lambda=\lambda^{\prime}$ such that $e^{-\lambda^{\prime}}=b^{\prime} / n$, the mean of $F_{\lambda^{\prime}}$ is by the lower bound calculation exactly $n \nu^{\prime}$, and we have chosen $v<\nu^{\prime}<a$. To conclude that $F_{\lambda^{\prime}}(n a)-F_{\lambda^{\prime}}(n v) \rightarrow 1$, instead of using a law of large numbers argument such as in the i.i.d. case, we simply compute the variance of $F_{\lambda^{\prime}}$ directly. Anticipating on the calculations of the next section, breaking the factor $e^{-\lambda x}$ in the Radon-Nikodym derivative of $F_{\lambda}$ into $e^{-\lambda \sum x_{i}}$ means that we can see $F_{\lambda}$ as a sum of independent Bernoulli random variables with parameter $\beta_{i}$ so that the variance is

$$
\operatorname{var} F_{\lambda^{\prime}}=\sum_{i=1}^{n} \beta_{i}\left(1-\beta_{i}\right) \leq \sum_{i=1}^{n} \beta_{i}=\operatorname{mean}\left(F_{\lambda}\right)=n \nu^{\prime}=O(n) .
$$

Another way to obtain this inequality is to do more direct computations:

$$
\begin{aligned}
\operatorname{var} F_{\lambda}= & \frac{\varphi_{n}^{\prime \prime}(\lambda)}{\varphi_{n}(\lambda)}-\left(\frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)}\right)^{2} \\
= & \sum_{i=1}^{n} \frac{e^{-\lambda}(1-1 / i)}{q_{i}}+\sum_{i \neq j} \frac{e^{-\lambda}(1-1 / i) e^{-\lambda}(1-1 / j)}{q_{i} q_{j}} \\
& -\left(\sum_{i=1}^{n} \frac{e^{-\lambda}(1-1 / i)}{q_{i}}\right)^{2} \\
= & \sum_{i=1}^{n} \frac{e^{-\lambda}(1-1 / i)}{q_{i}}-\sum_{i=1}^{n} \frac{e^{-2 \lambda}(1-1 / i)^{2}}{q_{i}^{2}} \\
\leq & \sum_{i=1}^{n} \frac{e^{-\lambda}(1-1 / i)}{q_{i}}=\frac{\varphi_{n}^{\prime}(\lambda)}{\varphi_{n}(\lambda)}=\operatorname{mean} F_{\lambda} .
\end{aligned}
$$

Since the variance is $O(n)$, by Chebyshev's inequality we have that $F_{\lambda}(n a)-$ $F_{\lambda}(n v) \rightarrow 1$. Therefore,

$$
\liminf _{n \rightarrow \infty} \frac{\log P\left[S_{n} \leq a n\right]}{n \log n} \geq v-1
$$

But $v$ is arbitrarily close to $a$, so the result is proved.
4. The uniform measure on $\boldsymbol{\partial} \boldsymbol{B}(\boldsymbol{a n})$. Let $\left\{\zeta_{i}^{\prime}, 1 \leq i \leq n\right\}$ have the distribution of $\left\{\zeta_{i}, 1 \leq i \leq n\right\}$ conditional on $\sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor$. Let $\left\{\zeta_{i}^{(\lambda)}, 1 \leq i \leq n\right\}$ be independent with distribution

$$
d F_{\lambda, i}(x)=\frac{1}{\phi_{i}(\lambda)} e^{-\lambda x} d F_{i}(x)
$$

where $F_{i}, \phi_{i}$ are respectively the distribution function and the Laplace transform of $\zeta_{i}$, and $\lambda$ is the optimal parameter of the previous section, $e^{-\lambda}=b / n$. It is easy to see that $\zeta_{i}^{(\lambda)}$ is another Bernoulli random variable with

$$
P\left[\zeta_{i}^{(\lambda)}=1\right]=P\left[\zeta_{i}=1\right] e^{-\lambda} \frac{1}{\phi_{i}(\lambda)}=\frac{1}{1+n /(b(i-1))}:=\beta_{i}
$$

We are now ready to prove:
THEOREM 3. Let $\left|\bigodot_{1}\right|$ be the length of the cycle that contains 1 . Under $\mu$, the uniform distribution on $\partial B(I$, an $)$,

$$
\left|\mathcal{C}_{1}\right| \Rightarrow \mathbf{G}
$$

where $\mathbf{G}$ is a geometric r.v. with $P(\mathbf{G}>k)=(b /(1+b))^{k}$ and $b$ satisfies $\log (1+$ b) $/ b=1-a$.

Sketch of the proof. The first part of demonstrating this is to recall what Arratia, Barbour and Tavaré [1] call the Feller coupling. Start with vertex 1 and choose $\sigma(1)$ uniformly from the $n$ possible choices. If this is 1 , then take vertex 2 and choose $\sigma(2)$ uniformly from the $n-1$ remaining possible choices. If $\sigma(1) \neq 1$, then choose $\sigma(\sigma(1))$ uniformly from the $n-1$ remaining choices, and so on, until the final vertex where there is only one possible choice. Although the construction is much different from the Chinese restaurant process, the reader should note that if $\xi_{i}$ is defined by $\xi_{i}=1$ if a cycle is not completed at the $i$ th stage and 0 otherwise, then $\left\{\xi_{i}: 1 \leq i \leq n\right\}$ and $\left\{\zeta_{i}: 1 \leq i \leq n\right\}$ have the same distribution.

From the last observation it follows that $N=\inf \left\{i: \xi_{i}=1\right\}$ has the same distribution as the length of the cycle containing 1 . We can now conclude the proof of the theorem, using the large deviation calculation of the volume, and an argument called the Gibbs conditioning principle (see [4]). This principle asserts that the distribution of the $\zeta_{i}$ conditional on $\sum_{i=1}^{n} \zeta_{i}=a n$ should be asymptotically
independent and their law given by that which minimizes the entropy, that is, the random variables $\zeta_{i}^{(\lambda)}$ with distribution

$$
\begin{equation*}
\frac{1}{\phi_{i}(\lambda)} e^{-\lambda x} d F_{i}(x) \tag{4.1}
\end{equation*}
$$

where $F_{i}, \phi_{i}$ are respectively the d.f. and the Laplace transform of $\zeta_{i}$, and $\lambda$ is the parameter that optimizes (3.1), that is, $e^{-\lambda}=b / n$.

Proof of Theorem 3. We will first need a lemma.
LEMMA 3. For any $n \geq 1$ and for every $\lambda>0$,

$$
\left(\zeta_{1}^{\prime}, \ldots, \zeta_{n}^{\prime}\right) \stackrel{d}{=}\left(\zeta_{1}^{(\lambda)}, \ldots, \zeta_{n}^{(\lambda)}\right) \quad \text { given } \sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor
$$

PROOF. Let $f_{1}, \ldots, f_{n}$ be bounded nonnegative Borel functions:

$$
\begin{aligned}
E\left(\prod_{i=1}^{n} f_{i}\left(\zeta_{i}^{\prime}\right)\right) & =E\left(f_{1}\left(\zeta_{1}\right) \cdots f_{n}\left(\zeta_{n}\right) \mid \sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor\right) \\
& =E\left(f_{1}\left(\zeta_{1}\right) \cdots f_{n}\left(\zeta_{n}\right) ; \sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor\right) P\left(\sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor\right)^{-1}
\end{aligned}
$$

On the other hand,

$$
\begin{aligned}
& E\left(f_{1}\left(\zeta_{1}^{(\lambda)}\right) \cdots f_{n}\left(\zeta_{n}^{(\lambda)}\right) ; \sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right) \\
& \quad=\int_{R^{n}} f_{1}\left(x_{1}\right) \cdots f_{n}\left(x_{n}\right) \mathbf{1}_{\left\{\sum x_{i}=\lfloor a n\rfloor\right\}} \prod_{i=1}^{n} d F_{i, \lambda}(x) \\
& \quad=\int_{R^{n}} \prod_{i=1}^{n} f_{i}\left(x_{i}\right) \mathbf{1}_{\left\{\sum_{i} x_{i}=\lfloor a n\rfloor\right\}} \prod_{i=1}^{n} \frac{e^{-\lambda x}}{\phi_{i}(\lambda)} d F_{i}\left(x_{i}\right) \\
& \quad=\frac{e^{-\lambda\lfloor a n\rfloor}}{\prod_{i=1}^{n} \phi_{i}(\lambda)} \int_{R^{n}} \prod_{i=1}^{n} f_{i}\left(x_{i}\right) \mathbf{1}_{\left\{\sum_{i} x_{i}=\lfloor a n\rfloor\right\}} \prod_{i=1}^{n} d F_{i}\left(x_{i}\right) \\
& \quad=\frac{e^{-\lambda\lfloor a n\rfloor}}{\prod_{i=1}^{n} \phi_{i}(\lambda)} E\left[f_{1}\left(\zeta_{1}\right) \cdots f_{n}\left(\zeta_{n}\right) ; \sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor\right]
\end{aligned}
$$

We can now divide and multiply by the probability of the events in the two sides of this equation to obtain that for some constant $C>0$

$$
E\left(\prod_{i=1}^{n} f_{i}\left(\zeta_{i}^{(\lambda)}\right) ; \mid \sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right)=C E\left(\prod_{i=1}^{n} f_{i}\left(\zeta_{i}\right) \mid \sum_{i=1}^{n} \zeta_{i}=\lfloor a n\rfloor\right)
$$

By taking $f_{1}=\cdots=f_{n}=1$ we see that $C=1$ and the lemma is proved.
We will need another lemma:
Lemma 4. The $\zeta_{i}^{(\lambda)}$ satisfy a local central limit theorem:

$$
P\left(\sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right) \sim C n^{-1 / 2}
$$

Proof. The proof of this local limit theorem follows very closely that of the usual i.i.d. case, which can be found in Theorem 5.2 of [5]. Let $\beta_{m}=P\left(\zeta_{m}^{(\lambda)}=1\right)$ [i.e., $\beta_{m}=(1+n / b(m-1))^{-1}$ ], and let $X_{m, n}=n^{-1 / 2}\left(\zeta_{m}^{(\lambda)}-\beta_{m}\right)$ be the rescaled Bernoulli variable. We start by noticing that $X_{m, n}$ satisfy the hypotheses of the Lindeberg-Feller theorem (Theorem 4.5 in [5]). Indeed, they are independent by definition; for all $\varepsilon>0, P\left(\left|X_{m, n}\right|>\varepsilon\right)=0$ as soon as $n^{-1 / 2} \leq \varepsilon$, since $\zeta_{m}^{(\lambda)} \leq 1$ and $\beta_{m} \leq 1$ as well. Moreover,

$$
\begin{aligned}
\sum_{m=1}^{n} E\left(X_{n, m}^{2}\right) & =\frac{1}{n} \sum_{m=1}^{n} \beta_{m}\left(1-\beta_{m}\right) \\
& \rightarrow \int_{0}^{1} \frac{x / b}{(1+(x / b))^{2}} d x:=\sigma^{2}
\end{aligned}
$$

Therefore $\sum_{m=1}^{n} X_{m, n} \Rightarrow \mathcal{N}(0, \sigma)$. At this point, the proof of the local limit theorem from [5] can be reproduced exactly. Therefore

$$
\sup _{x \in \mathbf{R}}\left|n^{1 / 2} P\left(\sum_{m=1}^{n} X_{m, n}=x\right)-n(x)\right| \rightarrow 0
$$

where $n(x):=\left(2 \pi \sigma^{2}\right)^{-1 / 2} \exp \left(-x^{2} / 2 \sigma^{2}\right)$. Since $\sum_{m=1}^{n} \beta_{m} \rightarrow a n$, and since $n(\cdot)$ is a continuous function, we can conclude the proof of the lemma by the above uniform convergence.

Now, by the Feller coupling, $\left|\mathcal{C}_{1}\right| \stackrel{d}{=} \inf \left\{k \geq 1: \zeta_{n-k}^{\prime}=0\right\}$, that is, we must reverse the time of the Chinese restaurant process. Hence by Lemmas 3 and 4:

$$
\begin{aligned}
P\left[\mathcal{C}_{1}>k\right] & =P\left[\zeta_{n}^{\prime}=1, \ldots, \zeta_{n-k+1}^{\prime}=1\right] \\
& =P\left[\zeta_{n}^{(\lambda)}=1, \ldots, \zeta_{n-k+1}^{(\lambda)}=1 \mid \sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right] \\
& =\frac{1}{P\left[\sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right]} P\left[\prod_{i=0}^{k-1} \zeta_{n-i}^{(\lambda)}=1 ; \sum_{i=1}^{n-k-1} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor-k\right]
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{1}{P\left[\sum_{i=1}^{n} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor\right]} \prod_{i=0}^{k-1} P\left[\zeta_{n-i}^{(\lambda)}=1\right] P\left[\sum_{i=1}^{n-k-1} \zeta_{i}^{(\lambda)}=\lfloor a n\rfloor-k\right] \\
& \sim \frac{1}{1+n /(b(n-1))} \cdots \frac{1}{1+n /(b(n-k))} \\
& \rightarrow \frac{1}{(1+1 / b)^{k}} .
\end{aligned}
$$

Hence Theorem 2 is proved.
5. Asymptotic singularity between $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$. In this section we give a proof of Theorem 5 that follows in an almost straightforward way from Theorems 2 and 3: $v$ and $\mu$ concentrate on permutations that have a different number of fixed points. First recall the statement of the theorem:

THEOREM 5. As $n \rightarrow \infty$, the hitting distribution $v$ and the uniform distribution $\mu$ on a sphere of radius an are asymptotically singular:

$$
d_{\mathrm{TV}}(\mu, \nu) \rightarrow 1
$$

LEMMA 5. The random partition of $\{1, \ldots, n\}$ derived from $v$ is exchangeable.

Proof. The probability to obtain a certain partition of $\{1, \ldots, n\}$ under $v$ only depends on the size of its blocs, which stays the same under the action of a given permutation. Hence $v$ yields an exchangeable partition of $\{1, \ldots, n\}$.

An immediate consequence is that the expected number of fixed points is $n v\left(\mathcal{C}_{1}=1\right)=n /(1+b)$. Next we show that under $v$ the number of fixed points $N$ is close to its expected value.

LEMMA 6.

$$
\operatorname{var} N=o\left(n^{2}\right)
$$

under $\nu$.
Proof. Let $x_{i}=\mathbf{1}_{\left\{\zeta_{i}^{\prime}=0 ; \zeta_{i+1}^{\prime}=0\right\}}$ be the indicator of the event that in the conditioned Chinese restaurant process, client number $i$ sits by himself. Then $N=\sum_{i} x_{i}$ and

$$
\begin{aligned}
\operatorname{var} N & =\sum_{i=1}^{n} \operatorname{var} x_{i}+\frac{1}{2} \sum_{i<j} \operatorname{cov}\left(x_{i}, x_{j}\right) \\
& \leq n+\frac{1}{2} \sum_{i<j} \operatorname{cov}\left(x_{i}, x_{j}\right)
\end{aligned}
$$

But when $j-i>1$, by the Gibbs asymptotic independence proved in Theorem 3, $\operatorname{cov}\left(x_{i}, x_{j}\right) \rightarrow 0$. Also, there are only $O(n)$ terms such that $j=i+1$ and in this case $\operatorname{cov}\left(x_{i}, x_{i+1}\right) \leq 1$, hence the sum $\sum_{i<j} \operatorname{cov}\left(x_{i}, x_{j}\right)=o\left(n^{2}\right)$.

To end the proof of Theorem 5 by Chebyshev's inequality there remains only to notice that:

Lemma 7. For $0<a<1$ and large enough $n$

$$
v\left(\left|\mathfrak{C}_{1}\right|=1\right) \neq \mu\left(\left|\mathfrak{C}_{1}\right|=1\right) .
$$

Proof. Recall that $b$ is defined by $\log (1+b) / b=1-a$. For $x \in(0,1)$, let $f(x)=1-\log (1+x) / x$, so that $b=f^{-1}(a)$.

On the other hand, an easy consequence of Berestycki and Durrett [3] or Theorem 0 is $\mu\left(\left|\mathfrak{C}_{1}\right|=1\right)=e^{-u^{-1}(a)}$. [Indeed, under $\mu,\left|\mathfrak{C}_{1}\right|$ is asymptotically the total progeny of a Poisson-Galton-Watson process, or $P G W$ process with offspring mean $u^{-1}(a)$.]

Hence the lemma is proved if we show that

$$
1 /(1+b) \neq e^{u^{-1}(a)} \quad \text { or } \quad u(x) \neq 1-x /\left(e^{x}-1\right)
$$

for all $x>0$.
We start by noticing that as $x \rightarrow 0, u(x) \sim x$ but $1-x /\left(e^{x}-1\right) \sim x / 2$. Hence $u(x)>1-x /\left(e^{x}-1\right)$ as $x \rightarrow 0$. The same is true as $x \rightarrow \infty$ [an easy argument shows indeed that $\left.u(x)=1-e^{-x}+o\left(e^{-x}\right)\right]$. Now those functions are both concave as we will see in a moment, hence this has to stay true on the whole open half-line $x>0$. (Notice that we have thus proved that the hitting distribution has always fewer fixed points than the uniform distribution.)

Lemma 8. The function $u$ appearing in Theorem 0 is concave.
Proof. For $c \leq 1$ this is obvious. When $c>1$, rather than carrying explicit calculations on the second derivative of $u$, we use a theoretic argument that exploits the recent result of Schramm [12], which says that the sizes of the pieces of the giant component in the random graph have approximately a Poisson-Dirichlet distribution. Since each fragmentation decreases the distance by 1 and each coalescence increases it by 1 , it is easy to see that

$$
\frac{d}{d c} \mathbf{E}\left[d\left(\sigma_{c n / 2}, I\right) \mid \mathcal{F}_{c n / 2}\right]=1-2 P\left[\text { fragm. } \mid \mathcal{F}_{c n / 2}\right]
$$

where $\mathcal{F}$. is the canonical filtration generated by the random walk. So we need to show that $P$ [fragm.] is an asymptotically increasing function of $c$. However, the probability of fragmenting a small cycle is asymptotically 0 (by duality and the fact
that $u$ is linear in the subcritical regime), and the probability of fragmenting one of the giant cycles can be computed explicitly using the Poisson-Dirichlet structure:

$$
P[\text { fragm. }] \rightarrow \mathbf{E} \sum_{i=1}^{\infty}\left(\theta(c) X_{i}\right)^{2}=\theta(c)^{2} \mathbf{E} \sum_{i=1}^{\infty} X_{i}^{2}=\frac{1}{2} \theta(c)^{2},
$$

where $\theta(c)$ is the survival probability of a $P G W(c)$ branching process and ( $X_{i}, i \geq 1$ ) follows the $P D(1)$ distribution. ( $\mathbf{E} \sum X_{i}^{2}=1 / 2$ follows from [10], formula (128).) Since $\theta(c)$ is an increasing function of $c$, the lemma is proved (and thus, so is Theorem 5).

REMARK. We have thus proved the formula

$$
u(c)=c / 2-\int_{0}^{c / 2} \theta(u)^{2} d u
$$

which is perhaps a little simpler to handle than the expression in Theorem 0.
6. Number of geodesics and Radon-Nikodym derivative. Here we prove the following theorem, which we will then use to prove a stronger version of the singularity theorem.

THEOREM 6. Suppose $d(\sigma)=t$ and $m_{1}, \ldots, m_{j}$ are the cycle lengths of $\sigma$. The number of paths of length $t$ from I to $\sigma$ is

$$
t!\prod_{i=1}^{j} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!} .
$$

From this it follows that if $t=[c n / 2]$ with $c<1$, then

$$
r(\sigma)=K_{n, t} \prod_{i=1}^{j} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!},
$$

where $K_{n, t}$ is a constant that only depends on $n$ and $t$.
Sketch of the proof. To see the first result, note that in order to go from $\sigma$ to $I$ in the shortest number of steps we must increase the number of cycles by 1 at each step, and to do this we must fragment a cycle at each step by transposing two of its elements. A cycle of length $m_{i}$ will require $m_{i}-1$ fragmentations. The first step in constructing a path is to decide on how to allocate the $t$ moves between the original cycles, which can be done in $t!/ \prod_{i=1}^{j}\left(m_{i}-1\right)$ ! ways. The next step is to count the number of ways that we can reduce a cycle of length $m_{i}$ in $m_{i}-1$ steps, which turns out to be simple: $m_{i}^{m_{i}-1}$.

Proof of Theorem 6. Given a partition of $\{1,2, \ldots, n\}$ into groups $A_{1}, \ldots, A_{j}$ of sizes $m_{i}, 1 \leq i \leq j$, the number of forests that consist of trees
with vertex sets $A_{1}, \ldots, A_{j}$ is by Cayley's formula for the number of unrooted trees on $m_{i}$ vertices

$$
\prod_{i=1}^{j} m_{i}^{m_{i}-2}
$$

Let $t=\sum_{i}\left(m_{i}-1\right)$. A given forest can be built up in $t$ ! ways so there are

$$
t!\prod_{i=1}^{j} m_{i}^{m_{i}-2}
$$

paths for our random graph process that end up producing a given partition. The number of permutations that correspond to a given partition is

$$
\prod_{i=1}^{j}\left(m_{i}-1\right)!.
$$

An equal number of paths end at each permutation with cycle sizes $m_{i}, 1 \leq i \leq j$, so the number of paths to a given permutation is

$$
t!\prod_{i=1}^{j} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!} .
$$

If $t=[c n / 2]$ with $c<1$, then the number of edge choices that end up producing no fragmentations is by Theorem 1 in [3]

$$
\sim\left(\binom{n}{2}\right)^{t} e^{-\kappa(c)}
$$

Taking the ratio of the last two results gives Theorem 6.
7. The size of the support of the hitting distribution. In this section we prove Theorem 7, restated below.

THEOREM 7. Suppose $a<1 / 2$. There exists a set $S_{n} \in \partial B(a n)$ such that $\nu\left(S_{n}\right) \rightarrow 1$ as $n \rightarrow \infty$ and

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \log \frac{\left|S_{n}\right|}{\mid \partial B(\text { an }) \mid}=\gamma<0 .
$$

Sketch of the proof. Obtaining a decay at least exponential is not very hard, even in the case $a>1 / 2$. However, it is not easy to prove that this is the correct rate for the decay of $|S| /|\partial B|$, and we restrict ourselves to the case $a<1 / 2$.

If $\sigma \in \partial B(a n)$, then we can use Theorem 6 to find that

$$
\log v_{0}(\sigma)=-a n \log n+a n+\sum_{k=1}^{n} a_{k} \log p_{k}+o(n)
$$

where $p_{k}$ is the Borel distribution with parameter $c$, and $a_{k}$ is the number of cycles of $\sigma$ of size $k$. But by the law of large numbers, $v_{0}\left(a_{k} / n\right)$ should have a limit as $n \rightarrow \infty$. Hence there is a set $S$ such that $\left(\log v_{0}(\sigma)+a n \log n\right) / n$ has a limit $-c_{1}$ whenever $\sigma \in S$. Because $v_{0}(S) \approx 1,|S| \approx \exp \left(a n \log n+c_{1} n\right)$. Moreover it is also true that $v(S) \rightarrow 1$. On the other hand, precise estimates on the size of $\partial B(a n)$ obtained via Kolchin's representation theorem tell us that $|\partial B(a n)|=\exp \left(a n \log n+c_{2} n+o(n)\right)$. (A statement of Kolchin's theorem can be found below.) Thus, the theorem holds with $\gamma=c_{1}-c_{2}$. To prove that $\gamma \neq 0$, we argue that the decay has to be at least exponential (a consequence of Kolchin's representation theorem).

Proof of Theorem 7. We will need precise estimates on the size of $\partial B(a n)$. Because we need estimates to order higher than just $n \log n$, sticking to the large deviation approach is not good enough. Rather, we will use Kolchin's representation theorem. We would like to thank Jim Pitman for pointing out this reference to us.

Suppose we can partition $\{1, \ldots, n\}$ into a certain number of clusters, which can all have different internal states. To be more specific, suppose that each partition of $\{1, \ldots, n\}$ into $k$ clusters leads to $v_{k}$ possible global states of the system $\{1, \ldots, n\}$, and that we can further assign each cluster of size $j$ one of $w_{j}$ possible internal states. We call such a combinatorial structure a ( $v, w$ )-partition (of $\{1, \ldots, n\})$. Kolchin's representation theorem answers with probabilistic means to the following purely combinatorial question: how many different $(v, w)$-partitions are there? Also, what does a random, uniform, $(v, w)$-partition look like?

Before going into the details of this theorem, let us see its relevance to our problem. The number of permutations at distance $a n$ from the identity is a special instance of the above Kolchin problem, where $v_{k}=\mathbf{1}_{\{k=(1-a) n\}}$ and $w_{j}=(j-1)!$. Indeed a permutation at distance $a n$ is exactly a permutation having $(1-a) n$ cycles and each cluster of size $j$ can be in one of the $(j-1)$ ! possible orderings of the cycle.

Here is the content of Kolchin's theorem (see [10]). Let $v(\theta)=\sum_{k=1}^{\infty} v_{k} \theta^{k} / k$ ! and let $w(\xi)=\sum_{j=1}^{\infty} w_{j} \xi^{j} / j$ ! be the so-called exponential generating function of the sequences $v$ and $w$. Let $K$ be an integer-valued random variable with distribution

$$
P(K=k)=v_{k} \frac{w(\xi)^{k}}{k!v(w(\xi))}
$$

and let $X$ be a random variable distributed according to

$$
P(X=j)=\frac{w_{j} \xi^{j}}{j!w(\xi)}
$$

Here $\xi$ is any parameter. In our setting, $K=(1-a) n$, a.s. and $X$ has the socalled logarithmic distribution, $P(X=j)=b^{j} / j \cdot \frac{1}{-\log (1-b)}$, for some parameter $b=w(\xi)$.

THEOREM 9 (Kolchin). The number of $(v, w)$-partitions is given by

$$
\frac{n!v(w(\xi))}{\xi^{n}} P\left(\sum_{i=1}^{K} X_{i}=n\right)
$$

where $X_{i}$ are i.i.d. samples of the variable $X$. Moreover, the sizes of the clusters in exchangeable random order have the same law as

$$
\left(X_{1}, \ldots, X_{K}\right) \quad \text { given } X_{1}+\cdots+X_{K}=n .
$$

For a precise definition of exchangeable random order, and further discussion of this theorem, see [10]. It is to be noted that here $\xi$ is any parameter. By playing on this parameter so as to make the event $S_{K}=n$ not unlikely (e.g., of probability $\propto n^{-1 / 2}$ rather than exponentially small), we get that the sizes of the clusters are approximately drawn from the r.v. $X$. Note that as a consequence we get here another proof of Theorem 3. Indeed, we see that the sizes of the cycles of a uniform permutation on $\partial B$ in exchangeable random order have a logarithmic distribution (asymptotically when the parameter is chosen suitably). Hence, a size-biased pick $\left|\mathcal{C}_{1}\right|$ should have distribution $P\left(X^{\prime}=j\right)=$ const. $j \cdot \frac{b^{j}}{j} \propto b^{j}$, a geometric random variable. The similarity between the large deviations-statistical mechanics approach and Kolchin's theorem is striking.

Another straightforward consequence of this theorem is the precise asymptotics for the size of a ball of radius $a n$. Indeed, in our setting, $v(\theta)=\theta^{(1-a) n} /[(1-a) n]$ ! and $w(\xi)=-\log (1-\xi)$, hence:

$$
|\partial B(a n)|=\frac{n!(-\log (1-\xi))^{(1-a) n}}{\xi^{n}((1-a) n)!} P\left(\sum_{i=1}^{(1-a) n} X_{i}=n\right)
$$

where $\xi$ is still any parameter. However, when $\xi$ is chosen such that $(1-a) \times$ $E(X)=1$, the local central limit theorem shows that $P\left(\sum_{i=1}^{(1-a) n} X_{i}=n\right) \sim$ $C n^{-1 / 2}$. By Stirling's formula, it is now straightforward to see that

$$
|\partial B(a n)|=\exp \left(a n \log n+c_{2} n+o(n)\right)
$$

Let us now turn our attention to the hitting distribution. We will get the corresponding estimate by analyzing the Radon-Nikodym derivative $r(\sigma)$ and the law of large numbers for $v$, as mentioned in the sketch of the proof.

More precisely, it follows from the proof of Theorem 6 that if $\sigma \in \partial B(a n)$, with cycle decomposition of size $m_{1}, \ldots, m_{(1-a) n}$, and $t=a n$, then

$$
v_{0}(\sigma)=\frac{1}{\binom{n}{2}^{t} e^{-\kappa(c)}} t!\prod_{i=1}^{n(1-a)} \frac{m_{i}^{m_{i}-2}}{\left(m_{i}-1\right)!}
$$

Let us write $a_{k}$ for the number of cycles of $\sigma$ of size $k$, so that $\sum_{k=1}^{n} a_{k}=n(1-a)$ and $\sum_{k=1}^{n} k a_{k}=n$. We can rewrite the above as

$$
\nu_{0}(\sigma)=\frac{1}{\binom{n}{2}^{t} e^{-\kappa(c)}} t!\prod_{k=1}^{n}\left(\frac{1}{c} \frac{k^{k-2}}{(k-1)!}\left(c e^{-c}\right)^{k}\right)^{a_{k}} c^{(1-a) n} \frac{1}{\left(c e^{-c}\right)^{n}}
$$

When we take the logarithm, calling $q_{k}=\frac{1}{c} \frac{k^{k-2}}{k!}\left(c e^{-c}\right)^{k}$ and $p_{k}=k q_{k}$,

$$
\begin{aligned}
\log v_{0}(\sigma)= & \kappa(c)-t \log \binom{n}{2}+\log t! \\
& +\sum_{k=1}^{n} a_{k} \log p_{k}+n(1-a) \log c-n \log c+c n
\end{aligned}
$$

Recalling that $t=a n, c=2 a$, and using Stirling's formula, we find that

$$
\begin{equation*}
\log v_{0}(\sigma)=-a n \log n+a n+\sum_{k=1}^{n} a_{k} \log p_{k}+o(n) \tag{7.1}
\end{equation*}
$$

We would now like to use the law of large numbers for $\nu_{0}$ since we know that under $v$,

$$
v\left(\left\{\sigma \in \partial B:\left|\frac{a_{k}}{n}-q_{k}\right| \leq \varepsilon ; \forall 1 \leq k \leq n\right\}\right) \rightarrow 1
$$

for given $\varepsilon>0$, and where $p_{k}$ is the Borel distribution. But it is not directly possible to take $S=\left\{\sigma \in \partial B:\left|\frac{a_{k}}{n}-q_{k}\right| \leq \varepsilon ; \forall 1 \leq k \leq n\right\}$ since we would obtain a bound $\varepsilon \sum_{k=1}^{\infty} \log p_{k}=\infty$. So we need to modify our choice: let
$S_{n}=\left\{\sigma \in \partial B:\left|\frac{a_{k}}{n}-q_{k}\right| \leq(\log n)^{-5} ; \forall 1 \leq k \leq(\log n)^{2}\right.$ and $a_{k}=0$ otherwise $\}$.
There are two things we need to check on $S_{n}$. First we need to see that it has the property that $v\left(S_{n}\right) \rightarrow 1$, and also that it has the correct size asymptotically. The first thing is taken care of by the next lemma, while the second will follow from the fact that $v_{0}\left(S_{n}\right) \rightarrow 1$, itself also a consequence of the lemma below.

Lemma 9. $v\left(S_{n}\right) \rightarrow 1$.
Proof. It suffices to prove that $v\left(\partial B-S_{n}\right) \rightarrow 0$. By the coupling with an Erdős-Renyi random graph, there is a $\beta>0$ such that no cycle can be greater than $\beta \log n$ with high probability under $v$. Hence it remains to prove that

$$
\sum_{k=1}^{(\log n)^{2}} v\left(\left|\frac{a_{k}}{n}-q_{k}\right| \leq(\log n)^{-5}\right) \rightarrow 0
$$

The basic idea is to use random graph estimates. Let $G(t)$ be the result of a random graph process where edges are added in a Poissonian way at rate $\binom{n}{2}$. Then the
expectation of the number of clusters of size $k a_{k}^{\prime}$ in $G(a n)$, is known to be $n q_{k}$ asymptotically with standard deviation $O\left(n^{1 / 2}\right)$, which is much smaller than the $n(\log n)^{-5}$ from the definition of $S_{n}$. So we need to show this still holds under $v$. Recall that we can couple the process ( $G(t), t \geq 0$ ) with a random walk ( $\sigma_{t}, t \geq 0$ ) where we multiply by a transposition $(i, j)$ whenever edge $(i, j)$ arrives in $G(t)$. Thus we may consider the first time $T$ that $\sigma$ is at distance $\lfloor a n\rfloor$ of the identity, and obviously a realization of $v$ is obtained as $\sigma_{T}$. We consider also $T^{\prime}$ the first time that $\lfloor a n\rfloor$ edges have been added to $G(t)$. Then since $T^{\prime}$ is the first time a Poisson process with intensity 1 exceeds the value $\lfloor a n\rfloor, T^{\prime}$ has mean $\lfloor a n\rfloor$ and variance $O(n)$. Hence the number of edges added to $G$ between $\lfloor a n\rfloor \wedge T^{\prime}$ and $\lfloor a n\rfloor \vee T^{\prime}$ is $O\left(n^{1 / 2}\right)$. On the other hand, only a bounded number $Z_{a n}$ of edges are added between $T^{\prime}$ and $T$, corresponding to the number of fragmentations at time $a n$, and thus all in all only $O\left(n^{1 / 2}\right)$ edges are added between $a n \wedge T$ and an $\vee T$, so this may create or destroy at most $O\left(n^{1 / 2}\right)$ clusters of size $k$ for each $1 \leq k \leq(\log n)^{2}$, which is much smaller than the $n(\log n)^{-5}$ from the definition of $S_{n}$. Hence we conclude that

$$
\sum_{k=1}^{(\log n)^{2}} v\left(\left|\frac{a_{k}}{n}-q_{k}\right| \leq(\log n)^{-5}\right) \rightarrow 0
$$

Thus $v\left(\partial B-S_{n}\right) \rightarrow 1$. Since $\nu_{0}$ is obtained as an asymptotically nondegenerate conditioning of $v$, it follows immediately that $\nu_{0}\left(\partial B-S_{n}\right) \rightarrow 0$ as well, that is, $\nu_{0}\left(S_{n}\right) \rightarrow 1$. We now show how to use this to estimate the size of $S_{n}$. By (7.1), for all $\sigma \in S_{n}$,

$$
\begin{aligned}
\frac{\log v_{0}(\sigma)+a n \log n}{n} & \geq a+\sum_{k=1}^{(\log n)^{2}}\left(q_{k}+(\log n)^{-5}\right) \log p_{k}+o(1) \\
& \geq a+\sum_{k=1}^{(\log n)^{2}} q_{k} \log p_{k}+o(1)
\end{aligned}
$$

from which we deduce that

$$
\liminf _{n \rightarrow \infty} \frac{1}{n}\left(\log v_{0}(\sigma)+a n \log n\right)=a+\sum_{k=1}^{\infty} q_{k} \log p_{k}:=-c_{1}
$$

After similar treatment for the lim sup, we get

$$
\frac{1}{n}\left(\log \nu_{0}(\sigma)+a n \log n\right) \rightarrow-c_{1}
$$

Since

$$
\nu_{0}\left(S_{n}\right)=\sum_{\sigma \in S_{n}} v_{0}(\sigma) \rightarrow 1
$$

it must be that $\left|S_{n}\right|=\exp \left(\right.$ an $\left.\log n+c_{1}+o(n)\right)$. Therefore

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \log \frac{\left|S_{n}\right|}{|\partial B|}=c_{1}-c_{2}:=\gamma
$$

It now remains to show that $\gamma \neq 0$. Observe that by Kolchin's theorem, we could pursue the asymptotic expansion of $|\partial B(a n)|$ and the next term would be polynomial in $n$. From the exact formula of $v(\sigma)$, we could also find the next term for $|S|$ and find that it is polynomial. Hence if $\gamma=0$, then we would have $|S| /|\partial B| \sim n^{-\alpha}$ for some $\alpha \geq 0$. But, another consequence of Kolchin's theorem is that the decay has to be at least exponential: for instance, permutations in $S$ have a number of fixed points characteristic of $v$ and not of $\mu$. As we have seen earlier the number of fixed points under $\mu, n /(1+b)$, is smaller than under the hitting distribution. But since the number of fixed points under $\mu$ is given by a sum of almost independent random variables,

$$
\sum_{i=1}^{n-\lfloor a n\rfloor} \mathbf{1}_{\left\{X_{i}=1\right\}} \quad \text { given } \sum_{i=1}^{n-\lfloor a n\rfloor} X_{i}=n
$$

we have that

$$
\begin{aligned}
\mu(S) & \leq P\left(\left.\frac{1}{n} \sum_{i=1}^{n-\lfloor a n\rfloor} \mathbf{1}_{\left\{X_{i}=1\right\}}<\frac{1}{1+b} \right\rvert\, \sum_{i=1}^{n-\lfloor a n\rfloor} X_{i}=n\right) \\
& \leq C n^{1 / 2} e^{-n \rho} \leq C^{\prime} e^{-n \rho^{\prime}}
\end{aligned}
$$

by standard large deviations (here, simply Markov's inequality), and because the event on which we condition is of probability $C n^{-1 / 2}$. Hence the decay has to be at least exponential and $\gamma$ cannot be 0 .

REMARK. The same argument shows that the hitting distribution of $\sigma$ is supported on a set at least exponentially small even in the case $a>1 / 2$, but of course we do not know whether this is a precise asymptotics. If the decay is still exponential after $a=1 / 2$, it seems likely that the exponential coefficient will not be smooth at $a=1 / 2$. In Figure 2 we have plotted the value of this coefficient against a time-change of $a$. It would be interesting to compute exact asymptotics in the case $a>1 / 2$ and make this picture complete.

REMARK. Kolchin's representation theorem could have been used already earlier for the proofs of Theorems 2 and 3 . This would actually simplify the proof of both results. However, we have chosen to keep the proofs as they were, because they do not rely on a technical result such as Kolchin's theorem, which is not as well known as standard large deviation theory.


FIG. 2. Numerical evaluation by Maple of the limiting behavior of $-n^{-1} \log \frac{\left|S_{n}\right|}{|\partial B|}$. The result is plotted as a function of $\xi=f^{-1}(a)$, where $f(\xi)=1+\log (1-\xi)(1-\xi) / \xi \cdot f^{-1}$ is an increasing function of $a$. This picture has been rigorously proved only for $a<1 / 2$, that is, $\xi<f^{-1}(1 / 2) \approx 0.715331863 \ldots$.
8. Asymptotic hyperbolicity under the uniform measure. Here we present a proof of Theorem 8. The sketch of the proof below contains some ideas that will be used and not re-explained in the actual proof that follows.

THEOREM 10. Let $0<a<1$ and let $\sigma$, $\pi$ be two random independent points chosen uniformly from $\partial B(a n)$. Then:

1. If $a<1-\log 2$,

$$
E(\sigma \mid \pi)_{p} \leq \delta(\log n)^{2}
$$

for some $0<\delta=\delta(a)<\infty$. Moreover, with probability asymptotically 1 , there is a geodesic between $\sigma$ and $\pi$ that comes within distance at most $\delta(\log n)^{2}$ of $p$.
2. If $a>1-\log 2$,

$$
E(\sigma \mid \pi)_{p} \sim \delta n
$$

for some $\delta=\delta(a)>0$. Moreover, no geodesic can approach $p$ closer than $\delta^{\prime} n$ for some $0<\delta^{\prime}<\infty$.

Sketch of the proof. To guess what the answer is, we exploit once again the connection with the theory of random graphs. The first thing to do is to realize that because of the symmetries of the Cayley graph $G_{n}$ it is enough to look at $d(I, \sigma \cdot \pi)$ and see whether it is approximately $2 a n$ or much smaller than $2 a n$.

To construct our graph, we will need some notation. Let

$$
\begin{equation*}
\pi=\tau_{1} \tau_{2} \cdots \tau_{a n} \tag{8.1}
\end{equation*}
$$

be a minimal decomposition of $\pi$ as a product of transpositions, with the following convention. If we list all cycles of $\pi$ in the order of their least element, then the transpositions $\tau_{i}$ are those $(x, y)$ such that $y$ comes just after $x$ in the cyclic decomposition of $\pi$ and $x$ and $y$ are in the same cycle, and we order the an transpositions according to their position in this canonical decomposition. To clarify the ideas, suppose

$$
\pi=(1437)(2)(58)(6109) ;
$$

then we write

$$
\pi=(14)(43)(37)(58)(610)(109) .
$$

We define the graph $\Gamma=(V, E)$ on $n(1-a)$ vertices as follows. Let $V=\{$ cycles of $\sigma\}$, and there is an edge between $\mathcal{C}$ and $\mathcal{C}^{\prime}$ if there is $x \in \mathcal{C}$ and $y \in \mathcal{C}^{\prime}$ such that $(x, y)$ is one of the an transpositions in the minimal decomposition described above. Note that this graph could have self-loops and multi-edges.

A notion that we will use on several occasions is that of being a terminal point. We say $x \in\{1, \ldots, n\}$ is terminal if $x$ does not appear more than once in the transpositions of the above minimal decomposition. This means that, with those conventions, $x$ is situated at the "end" of the cycle of $\pi$ in which it is contained.

Here is why we are interested in the properties of the graph $\Gamma$. If we define $\sigma_{0}=\sigma$ and, for $1 \leq r \leq a n, \sigma_{r}=\sigma \cdot \tau_{1} \cdots \tau_{r}$, and consider the process ( $\sigma_{r}, 0 \leq r \leq a n$ ), this is a walk on $G_{n}$ starting at $\sigma$ and ending at $\sigma \cdot \pi$. Moreover, since at each step we are multiplying by a transposition, the cycles of $\sigma_{r}$ evolve according to a discrete coagulation-fragmentation chain, with cycles merging when
the transposition involves elements from different cycles, and cycles splitting otherwise (as it is the case for simple random walk on $G_{n}$ ). Therefore, $\Gamma$ is the graph that results from drawing an edge between two cycles of $\sigma$ as we encounter a transposition joining those two cycles. In particular, the same argument that shows that the Erdős-Renyi random graph is an upper bound for the sizes of the cycles of simple random walk on $G_{n}$, will show that the cycles of $\sigma \cdot \pi$ are subcomponents of the connected components of $\Gamma$, with possibility of fragmentation whenever there is a cycle in $\Gamma$, or a self-loop or a multiple-edge. All other edges represent coalescence of cycles in the walk ( $\sigma_{r}, 0 \leq r \leq a n$ ).

In particular, the property of $\Gamma$ that we will be most interested in, will be to decide whether $\Gamma$ has a giant component, meaning a component containing a positive fraction of all $n(1-a)$ vertices. Indeed when all cycles of $\Gamma$ are small, we should expect very few cycles in $\Gamma$ and hence little fragmentation. Hence most steps of the walk $\sigma_{r}$ are coalescence events and the number of cycles decreases linearly; in other words, in the case that all cycles are small, $d(\sigma, \pi) \approx 2 a n$. On the other hand, if $\Gamma$ contains a giant cycle, then we can expect many cycles in the graph and hence many fragmentation events in the walk ( $\sigma_{r}, 0 \leq r \leq a n$ ), which means that $d(\sigma, \pi) \ll 2 a n$.

Here is our strategy to see whether there is a giant component in $\Gamma$. Rather than counting the number of cycles of $\sigma$ that a component of $\Gamma$ contains, we prefer to compute the exact number of integers in $\{1, \ldots, n\}$ that it actually encloses prior to shrinking all cycles of $\sigma$ into points. Formally, this means, give weight $W(\mathcal{C})=|\mathcal{C}|$ to any vertex $\mathcal{C}$ of $\Gamma$, and ask what is the total weight of a connected component of $\Gamma$. Let $\mathcal{C}_{1}(\Gamma)$ denote a size-biased pick from the connected components of $\Gamma$, that is, the total weight of the connected component of $\Gamma$ containing " 1 " [or, more precisely, $\left.\mathcal{C}_{1}(\sigma)\right]$.

Lemma 10 shows that $W\left(\mathcal{C}_{1}(\Gamma)\right)$ converges in distribution to the total progeny of a branching process with offspring distribution a shifted geometric random variable. The idea is that by Theorem 3, the various cycles of $\sigma$ are asymptotically i.i.d., so that each edge in $\Gamma$ adds to the weight of $\mathcal{C}_{1}(\Gamma)$ a contribution which is, by Theorem 3 , asymptotically a geometric random variable $\mathbf{G}$ with parameter $1 /(1+b)$ where $b$ satisfies $\log (1+b) / b=1-a$. This seems to give an infinite progeny almost surely (since $\mathbf{G} \geq 1$ a.s.). However, to every point that we examine there is a positive probability that it is a terminal point. In this case, that integer does not connect to a new independent cycle of $\sigma$, and hence its offspring is 0 . This kind of modified branching process is defined more precisely and analyzed in Section 8.3. The key fact is that because of the special properties of the asymptotic law of $\pi$, which involves geometric random variables, this modified branching process is in fact equal in distribution to another branching process where the offspring distribution has been shifted from $\mathbf{G}$ to $\mathbf{G}-1$. In all that follows, we call $X$ a random variable such that

$$
X \stackrel{d}{=} \mathbf{G}-1
$$

Hence $\Gamma$ has a giant component if, and only if, $E(\mathbf{G})>2$. Since $p=1 /(1+b)$ and $\log (1+b) / b=1-a$,

$$
P(\Gamma \text { has a giant component })>0 \Longleftrightarrow a>1-\log 2 .
$$

## Proof of Theorem 8.

8.1. Structure of the proof. As this proof is rather long, we feel that it is appropriate to explain how the various arguments are used. In Section 8.2, we prove that $W\left(\mathcal{C}_{1}(\Gamma)\right) \Rightarrow \sum_{t \geq 0} Z_{t}$ the total progeny of a branching process with offspring distribution $X$. Then in Section 8.3, we define a modified branching process, and prove that in the case of geometric random variables this becomes another branching process with shifted offspring distribution. We then use this to prove by hand that in the subcritical case, $\left|\mathcal{C}_{1}(\sigma \cdot \pi)\right|$ is dominated by such a modified branching process. Since this is a subcritical branching process, we prove the exponential decay of the tail of $\left|\mathcal{C}_{1}(\sigma \cdot \pi)\right|$, uniformly in $n$. This enables us to show that as long as $a<1-\log 2$ there are very few fragmentation events in the walk $\left(\sigma_{r}, r=0, \ldots, a n\right)$. The supercritical case is treated in Section 8.5. Since we have established branching process asymptotics, we can use the duality principle of a branching process between the subcritical phase and the supercritical phase. This shows that the number of clusters of $\Gamma$ in the supercritical regime can be computed by looking at the number of clusters of $\Gamma$ for some specific subcritical time. Since we have proved that the distance is linear in this regime, we now know how many clusters $\Gamma$ has at any subcritical time, and it follows that the number of clusters of $\Gamma$ in the supercritical regime is strictly less than what it would be if the distance was still linear. It only remains to prove that at any given time the number of extra cycles that were generated by some fragmentation (and have not been reabsorbed by other large cycles) is $O\left(n^{1 / 2}\right)$, which is done in Section 8.6.
8.2. Branching process asymptotics. To start proving things, we need some more notation. Let $A_{0}^{n}=\{1\}$ and define recursively the $A_{k}^{n}$ by

$$
A_{k+1}^{n}=\bigcup_{x \in A_{k}^{n}}\left\{\mathcal{C}_{\pi(x)}(\sigma)\right\}-\bigcup_{1 \leq j \leq k} A_{j}^{n} .
$$

The $A_{k}^{n}$ correspond to growing the branching process generation after generation, rather than cycle after cycle. Let $\left(Z_{t}, t=0,1, \ldots\right)$ be a branching process with offspring distributed as $X$. Note that by the construction of $\Gamma$, we also have that

$$
\sum_{k=0}^{\infty}\left|A_{k}^{n}\right|=W\left(\bigodot_{1}(\Gamma)\right)
$$

Lemma 10. As $n \rightarrow \infty$,

$$
\left(\left|A_{0}^{n}\right|,\left|A_{1}^{n}\right|, \ldots\right) \Rightarrow\left(Z_{0}, Z_{1}, \ldots\right)
$$

Proof. Let us start by the convergence of $\left(\left|A_{0}^{n}\right|,\left|A_{1}^{n}\right|\right)$. If $j=0, P\left(\left|A_{0}^{n}\right|=1\right.$, $\left.\left|A_{1}^{n}\right|=0\right)=P(\pi(1)=1) \rightarrow 1 /(1+b):=p=P(X=0)$. If $j \geq 1$, then

$$
\begin{aligned}
P\left(\left|A_{0}^{n}\right|=1,\left|A_{1}^{n}\right|=j\right) & =P\left(\pi(1) \neq 1 ;\left|\mathcal{C}_{\pi(1)}(\sigma)\right|=j\right) \\
& =P(\pi(1) \neq 1) \cdot P\left(\left|\mathcal{C}_{\pi(1)}(\sigma)\right|=j \mid \pi(1) \neq 1\right) \\
& \rightarrow(1-p) \cdot(1-p)^{j-1} p=P(X=j)
\end{aligned}
$$

Indeed, conditionally on $\{\pi(1) \neq 1\}, \pi(1)$ is uniform on $\{2, \ldots, n\}$, so that $\mathcal{C}_{\pi(1)}(\sigma)$ is as good a size-biased pick as $\mathcal{C}_{1}(\sigma)$, and we can apply Theorem 3.

Now let us consider the general case of finite-dimensional distributions. Let $n_{1}>0, n_{2}>0, \ldots, n_{k} \geq 0$ with $\sum_{i} n_{i} \leq n$. We are trying to compute the asymptotics of

$$
P\left(\left|A_{0}^{n}\right|=1,\left|A_{1}^{n}\right|=n_{1}, \ldots,\left|A_{k}^{n}\right|=n_{k}\right) .
$$

To do this, we need to evaluate the probability of a collision occurring in the first $k$ stages, that is,

$$
P\left(\pi(x) \in \bigcup_{1 \leq i \leq k} A_{i}^{n}-\mathcal{C}_{x}(\pi) \text { for some } x \in A_{j}^{n} \text { with } j \leq k\right)
$$

We will say of an $x$ such as in the event above, that it makes a backward connection. Hence an $x$ makes a backward connection if $\pi(x)$ maps it to some lower level in the branching process, but $x$ is not a terminal point. Therefore backward connections (or collisions) are exactly those that may lead to a fragmentation, as explained in the sketch of the proof.

It is easy to see that $P$ (collision in first $k$ stages) $=O(1 / n)$. In fact, it follows from the uniformity of Lemma 11 that

$$
\begin{aligned}
P\left(\left|A_{0}^{n}\right|=n_{0}, \ldots,\left|A_{k}^{n}\right|=n_{k} ; \text { b.w. collisions in } k \text { first stages }\right) & \leq \sum_{j=1}^{k} n_{j} \frac{\sum_{i=0}^{j-1} n_{i}}{n} \\
& \leq \frac{\left(\sum_{i=0}^{k} n_{i}\right)^{2}}{n}
\end{aligned}
$$

(see also Lemma 12 where similar estimates are derived).
Therefore it is enough to consider

$$
P\left(\left|A_{0}^{n}\right|=1,\left|A_{1}^{n}\right|=n_{1}, \ldots,\left|A_{k}^{n}\right|=n_{k} \mid\right. \text { given no b.w. connection). }
$$

Suppose $A_{k-1}^{n}=\left\{x_{1}, \ldots, x_{n_{k-1}}\right\}$. Conditionally on the event that there is no collision in the $k$ first stages, $\pi\left(x_{1}\right), \ldots, \pi\left(x_{n_{k-1}}\right)$ belong to yet unexplored cycles (as long as they are not terminal). After decomposition on the number of such $x$
[call $\mathrm{T}\left(A_{k-1}^{n}\right)$ the number of terminal elements in the set $A_{k-1}^{n}$ ], the last probability is equal to

$$
\begin{aligned}
& =\sum_{j=1}^{n} P\left(\sum_{i=1}^{n_{k-1}-j}\left|\mathcal{C}_{\pi\left(x_{i}\right)}(\sigma)\right|=n_{k}\right) P\left(\mathrm{~T}\left(A_{k-1}^{n}-j \mid \text { no b.w. conn. }\right)\right) \\
& \rightarrow \sum_{j \geq 0} P\left(\sum_{i=1}^{n_{k-1}-j} X_{i}=n_{k}\right) p^{j} \\
& =P\left(\sum_{i=1}^{n_{k-1}} X_{i}=n_{k}\right)
\end{aligned}
$$

by the asymptotic independence property of a finite number of size-biased cycles, and the fact that given there was no backward connection, the $x$ in level $A_{k-1}^{n}$ belong to different cycles of $\pi$, so that the events that they are terminal are independent asymptotically.

These are the transition probabilities of a branching process with offspring distribution $X$, so the lemma is proved.
8.3. A modified branching process. One way to formalize the idea that a vertex has a geometric number of children only during finitely many generations, is to use a modified branching process where each individual $x$ is endowed with a nonnegative, integer-valued random variable $T(x)$, that represents the "life-time" of its family. As long as $T(x)>0, x$ will keep having children according to the original offspring progeny $L$. But when $T(x)=0$, the individual will be declared "terminal" and will not be allowed to have any children.

Here is a rigorous description of this modified branching process. Let $X_{t, i}$ be a collection of i.i.d. random variables with distribution $L$, a fixed distribution on the nonnegative integers (the original progeny). Let $T_{t, i}$ be i.i.d. nonnegative integervalued random variables, distributed according to another distribution $L^{\prime}$, the lifetime. Let $Z_{t}$ be the size of the process at time $t$ (with discrete time). Define $Z_{0}=1$, and give the root lifetime $T_{0,0}$. Then define recursively $Z_{t}$ by

$$
\begin{equation*}
Z_{t+1}=\sum_{i=0}^{Z_{t}} X_{t, i} \mathbf{1}_{\left\{T\left(x_{i}\right)>0\right\}} \tag{8.2}
\end{equation*}
$$

where $x_{1}, \ldots, x_{Z_{t}}$ are the $Z_{t}$ individuals of generation $t$. If $y_{1}, \ldots, y_{Z_{t+1}}$ are the $Z_{t+1}$ individuals of generation $t+1$, the rule that we adopt for the value of $T\left(y_{1}\right), \ldots, T\left(y_{Z_{t+1}}\right)$ is the following. If $T\left(x_{i}\right)>0$, give all $X_{t, i}$ children of $x_{i}$ independent lifetimes from $T_{t, i}$, except for one of its children, say $y_{j}$, for which $T\left(y_{j}\right):=T\left(x_{i}\right)-1$. Rigorously, let $N_{t}=\#\left\{i: T\left(x_{i}\right)>0\right\}$, rewrite the $x_{i}$ 's removing the terminal ones and call them $x_{1}^{\prime}, \ldots, x_{N_{t}}^{\prime}$. Let $T\left(y_{1}\right)=T\left(x_{1}^{\prime}\right)-$
$1, \ldots, T\left(y_{N_{t}}\right)=T\left(x_{N_{t}}^{\prime}\right)-1$, and let $T\left(y_{N_{t}+1}\right)=T_{t+1, N_{t}+1}, \ldots, T\left(y_{Z_{t+1}}\right)=$ $T_{t+1, Z_{t+1}}$.

Of course we make this definition because asymptotically, $W\left(\mathcal{C}_{1}(\Gamma)\right)$ will be well approximated by such a system, where the offspring $L$ is the size of a cycle of $\sigma$, and where $L^{\prime}$ is the size of a cycle of $\pi$. Indeed, suppose we are exploring the cluster containing 1 in the graph of the superposition of $\sigma$ and $\pi . T(1)$ is $\left|\mathcal{C}_{1}(\pi)\right|$, which corresponds to the fact that after that many iterations of $\pi$ we are back to where we started and no longer add anything new to the population of the cluster. However, after one iteration say, all vertices in the first generation, other than $\pi(1)$ itself, belong to different cycles of $\pi$ with high probability. Therefore their lifetime should be an independent random variable, distributed as $L^{\prime}$.

In general, the ageing branching process $\left(Z_{t}, t=0,1, \ldots\right)$, where each individual has a "lifetime" that it transmits to one of its children, is not a Markov process with respect to its own filtration $\sigma\left(Z_{0}, Z_{1}, \ldots\right)$. Indeed the size of the generation $t+1$ depends not only on the size of generation $t$, but also on the random variables $T_{x}$ where $x$ is an individual of generation $t$, so one would need to add in the filtration the values of $T(x)$ for each generation.

However, a miracle happens due to the fact that the cycles of $\pi$ have (asymptotically) a geometric distribution $\mathbf{G}$. Let $p^{\prime}$ be the parameter of $\mathbf{G}: P(\mathbf{G}=j)=$ $\left(1-p^{\prime}\right)^{j-1} p^{\prime}$. Then the distribution $L^{\prime}$ of the random variables $T_{t, i}$ is again $\mathbf{G}$. For $k \geq 1$, conditionally on $T>k, T-k$ is distributed as $\mathbf{G}$. This fact, called "lack of memory," has the following amazing consequence:

Proposition 1. When the lifetime $L^{\prime}$ is a geometric random variable, $\left(Z_{t}, t=0, \ldots\right)$ is a Markovian branching process with offspring distribution $L \mathbf{1}_{\left\{L^{\prime}>0\right\}}$. When $L \stackrel{d}{=} L^{\prime} \stackrel{d}{=} \mathbf{G}$, this distribution is $X \stackrel{d}{=} \mathbf{G}-1$.

Proof. Let $B_{t, i}$ be Bernoulli random variables with success parameter $P\left(B_{t, i}=1\right)=p^{\prime}$. Because $\mathbf{G}={ }_{d} \inf \left\{t \geq 0 ; B_{t, i}=1\right\}$, the event $\left\{T\left(x_{i}\right)>0\right\}$ is the same as $\left\{B_{t, i}=0\right\}$, so (8.2) becomes

$$
\begin{equation*}
Z_{t+1}=\sum_{i=0}^{Z_{t}} X_{t, i} \mathbf{1}_{\left\{B_{t, i}=0\right\}} \tag{8.3}
\end{equation*}
$$

This expresses the fact that for each new vertex visited, we can take the decision of closing the cycle, independently of the past. When the cycle still has some length to be explored, then the vertex has $X_{t, i}$ children. This decision affects the law of progeny at a given vertex. The new distribution of the progeny is now, by (8.3):

$$
P(X=j)= \begin{cases}p^{\prime}, & \text { if } j=0  \tag{8.4}\\ \left(1-p^{\prime}\right) P\left(X_{t, i}=j\right), & \text { if } j \geq 1\end{cases}
$$

Of course, for our problem, $\sigma={ }_{d} \pi$, so both $L$ and $L^{\prime}$ are distributed as $\mathbf{G}$. As can be readily checked from (8.4), the distribution of $X$ is thus a shift of $\mathbf{G}$ :

$$
\begin{equation*}
X \stackrel{d}{=} \mathbf{G}-1 . \tag{8.5}
\end{equation*}
$$

8.4. Fragmentations in the subcritical case. First note that if $\sigma$ is a uniform permutation on $\partial B(a n)$, if we visit all points in $\{1, \ldots, n\}$ according to their order of appearance in the canonical cyclic decomposition of $\sigma$, and call this process $V_{t}(0 \leq t \leq n)$, then the successive points are in some sense uniformly chosen from what remains to be found, at least as long as we do not have to start a new cycle. More precisely:

LEMMA 11. Given $\left(V_{0}, \ldots, V_{t}\right)$ and given that $V_{t}$ is not a terminal point,

$$
\sigma\left(V_{t}\right) \text { is uniform on }\{1, \ldots, n\}-\left\{V_{0}, \ldots, V_{t}\right\} .
$$

The proof of this lemma follows directly from the Feller coupling presentation of a uniform permutation on $G_{n}$, which also has (obviously) this property. Conditioning on the number of cycles does not change how the cycles are filled in.

Lemma 12. Suppose the branching process is subcritical, that is, $p>1 / 2$ or (equivalently) $a<1-\log 2$. Then the number of fragmentations in the walk $\left(\sigma_{r}, r=0, \ldots, n\right)$ is $o(n)$.

Basically, all cycles are fairly small, so by improving our estimates on the number of collisions, we should get an $O(1)$ bound, just like in the Erdős-Renyi case. Technicality arises due to the fact that the cycles are conditioned independent random variables, and not just independent. Here is a rigorous proof.

Proof of Lemma 12. We prove things in two steps.
First, we prove a uniform bound for the size of a cluster: we show that if we write $\pi=\prod_{i=1}^{a n} \tau_{i}$, and denote $\pi_{r}:=\prod_{i=1}^{r} \tau_{i}$,

$$
\begin{equation*}
P\left(\left|\mathbb{C}_{1}\left(\sigma \cdot \pi_{r}\right)\right|>u \text { for some } r \leq a n\right) \leq C n \exp (-\alpha u), \tag{8.6}
\end{equation*}
$$

where $C$ and $\alpha$ are constants independent of $n$, and $u$ is any number.
Once this exponential control is proved, we can bound the number of times that one of the $\tau_{r}$ 's will yield a fragmentation. Indeed, recall that to obtain $\sigma \cdot \pi$ we can perform successively the $\tau_{r}$ 's on $\sigma$, and each one yields a coagulation or a fragmentation. We hence view this as a process indexed by $1 \leq r \leq a n$. In the course of this process, at all times, by (8.6) applied to $u=(\log n)^{2}$, no cluster is larger than $(\log n)^{2}$ with overwhelming probability, so that by Lemma 11:

$$
P\left(\tau_{r+1} \text { yields a fragmentation }\right) \leq 2(\log n)^{2} / n
$$

There are (exactly) an transpositions to perform, hence:

$$
E(\# \text { frag. }) \leq 2 a(\log n)^{2}
$$

This is already largely enough to prove Lemma 12.

We will now prove that (8.6) holds, since this is the only thing that remains to be proved. Although we have seen that in the limit each cluster is a subcritical branching process (for which such an exponential tail of the total progeny holds), when $n$ is finite there is no real branching process available to dominate $\mathcal{C}_{1}\left(\sigma \cdot \pi_{r}\right)$, essentially because the sizes of the cycles are not i.i.d. random variables. However, they are conditionally independent (cf. Theorem 9, or Theorem 3), and we will use this fact to construct a real branching process that dominates $\mathcal{C}_{1}\left(\sigma \cdot \pi_{r}\right)$, when conditioned on some mild event. This conditioning accounts for the extra factor $n$ in (8.6).

Here is how we proceed. By Kolchin's representation theorem (Theorem 9), there are random variables $\left(X_{1}, \ldots, X_{n(1-a)}\right)$ such that the joint law of the sizes of the cycles of $\sigma$ is $\left(X_{1}, \ldots, X_{n(1-a)}\right)$ given $\sum_{i} X_{i}=n$ (we will call $A_{n}$ the event that $\sum_{i} X_{i}=n$ ). The $X_{i}$ 's constitute a "pool" of possible cycle sizes. Similarly, there are random variables $\left(Y_{1}, \ldots, Y_{n(1-a)}\right)$ such that the joint law of the sizes of the cycles of $\pi$ is $\left(Y_{1}, \ldots, Y_{n(1-a)}\right)$ given $\sum_{i} Y_{i}=n$ (let $B_{n}$ be the event that $\sum_{i} Y_{i}=n$ ).

We give an upper bound of $\mathcal{C}_{1}(\sigma \cdot \pi)$ in terms of the modified branching processes of Section 8.3, that uses only the $X_{i}$ 's and the $Y_{i}$ 's. Start with vertex 1 and choose a size-biased pick $X_{1}^{\prime}$ of the $X_{i}$ 's (the cycle containing 1). Put $T(1)=Y_{1}^{\prime}$, a size-biased pick of the $Y_{i}$ 's. Next, given $X_{1}^{\prime}=k$, put $T(2)=$ $Y_{2}^{\prime}, \ldots, T(k)=Y_{k}^{\prime}$. All vertices with positive lifetime $T$ have a number of children given by a size-biased pick of the remaining $X_{i}$ 's. They transmit their lifetime -1 to one of their children and the rest have lifetimes given by size-biased picks from the remaining $Y_{i}$ 's. Then repeat the procedure until we cannot go any further (i.e., until all vertices at a given generation have lifetime $T=0$, or until all $X_{i}$ 's and $Y_{i}$ 's have been picked). Call $Z^{\prime}$ the total population obtained at the end of this construction.

We claim that $Z^{\prime}$ dominates all stages of $\mathcal{C}_{1}\left(\sigma \cdot \pi_{r}\right)$, because $Z^{\prime}$ gives the cycles of $\sigma$ coagulated by those of $\pi$, without taking any account of eventual fragmentations. In particular, in $Z^{\prime}$, as long as a vertex $x$ is not terminal $(T(x)>0)$, the children of $x$ will be part of the population of $Z^{\prime}$. Of course in the event of a collision or a backward connection, $Z$ does not contain any additional children, so that $Z<Z^{\prime}$. Therefore

$$
\begin{aligned}
P(Z>u) & \leq P\left(Z^{\prime}>u \mid A_{n} \text { and } B_{n}\right) \\
& \leq P\left(A_{n}\right)^{-1} P\left(B_{n}\right)^{-1} P\left(Z^{\prime}>u ; A_{n} ; B_{n}\right) \\
& \leq \operatorname{Cn} P\left(Z^{\prime}>u\right) .
\end{aligned}
$$

Indeed, by the local central limit theorem (see [5]),

$$
P\left(A_{n}\right)=P\left(B_{n}\right)=P\left(\sum_{i=1}^{n(1-a)} X_{i}=n\right) \sim C n^{-1 / 2}
$$

To complete the proof, it remains to notice that size-biasing the logarithmic distribution of Kolchin's theorem gives a geometric random variable. Therefore, by arguments already developed in the sketch of the proof, $Z^{\prime}$ is the total population of a branching process with offspring distribution $X$ of (8.5), and starting with a geometric number of individuals $\mathbf{G}$. Because $p>1 / 2$, this branching process is subcritical. In this case, classical estimates [2,5] show the exponential tail

$$
P\left(Z^{\prime}>u\right) \leq C \exp (-\alpha u)
$$

This concludes the proof of (8.6), and also that of Lemma 12.
REMARK. It is possible to avoid the use of Kolchin's representation theorem in the above proof. Indeed, by Theorem 3, a size-biased pick of the cycles has, after unconditioning on some event of probability $\propto n^{-1 / 2}$, a distribution which is given by the lengths of sequences of 1 in the Bernoulli trials $\zeta_{i}^{(\lambda)}$. However, since $\beta_{i}=P\left(\zeta_{i}^{(\lambda)}=1\right) \leq b /(1+b)$, it follows that the distribution of a size-biased cycle is thus (after unconditioning) stochastically dominated by the geometric random variable $\mathbf{G}$.
8.5. Mean in the supercritical regime and duality. Although this may seem a little surprising at first, we use the result from the subcritical case to get that for the supercritical case. The idea is to use the duality of branching processes.

A crucial remark is that the number of cycles of $\sigma \cdot \pi$ is given by $\sum_{x=1}^{n} 1 /$ $\left|\mathcal{C}_{x}(\sigma \cdot \pi)\right|$, hence by exchangeability

$$
\frac{1}{n} E(\# \text { clusters of } \Gamma)=E\left(\frac{1}{\left|\mathcal{C}_{1}(\Gamma)\right|}\right) \rightarrow E\left(\left.\frac{1}{T-1} \right\rvert\, T>1\right),
$$

where $T$ is the total progeny of a branching process with offspring distributed as $X$ and started with one individual. Indeed, let us not forget that the first generation $A_{0}$ of the branching process is itself a geometric random variable, so we can add an imaginary root and then subtract it (thus $T-1$ ). Introducing an extra vertex for the root allows us to make use of the duality principle of branching processes [2, 6].

The duality principle states that a supercritical branching process, conditioned on extinction, is another branching process, subcritical, whose offspring distribution is given through its generating function. If $\phi(s)=E\left(s^{X}\right)$ is the generating function of $X$ and $\alpha$ is the extinction probability $\alpha=P(T<\infty)$, then the conditioned process has offspring distribution characterized by

$$
\phi^{\prime}(s)=\phi(s \alpha) / \alpha
$$

Here, $P(X=j)=(1-p)^{j} p$, so

$$
\phi(s)=\frac{p}{1-s(1-p)}
$$

Therefore

$$
\phi^{\prime}(s)=\frac{p / \alpha}{1-s \alpha(1-p)} .
$$

The fixed point equation for $\alpha$ yields that

$$
\begin{equation*}
\frac{p}{\alpha}+\alpha(1-p)=1 \tag{8.7}
\end{equation*}
$$

so that $\phi^{\prime}$ is the Laplace transform of another shifted geometric random variable $X^{\prime}$, with parameter $p^{\prime}=p / \alpha$. Let $T^{\prime}$ be the total progeny of a branching process with offspring $X^{\prime}$ and started with one individual.

Let us now relate the supercritical and subcritical regimes. By duality,

$$
\begin{aligned}
E\left(\left.\frac{1}{T-1} \right\rvert\, T>1\right) & =\frac{P(T<\infty)}{P(T>1)} E\left(\left.\frac{1}{T-1} \mathbf{1}_{\{T>1\}} \right\rvert\, T<\infty\right) \\
& =\frac{P(T<\infty)}{P(T>1)} E\left(\frac{1}{T^{\prime}-1} \mathbf{1}_{\left\{T^{\prime}>1\right\}}\right) \\
& =\frac{P(T<\infty)}{P(T>1)} P\left(T^{\prime}>1\right) E\left(\left.\frac{1}{T^{\prime}-1} \right\rvert\, T^{\prime}>1\right)
\end{aligned}
$$

However, for the subcritical regime, we know by Lemma 12 that there are only $o(n)$ fragmentations, so the distance between $\sigma$ and $\pi$ is $2 a n+o(n)$ and the number of clusters is $(1-2 a) n+o(n)$. Hence $E\left(\left.\frac{1}{T^{\prime}-1} \right\rvert\, T^{\prime}>1\right)=1-2 a^{\prime}$, where $a^{\prime}$ is the radius corresponding to the conditioned parameter $p^{\prime}$. Since $p=1 /(1+b)$ and $\log (1+b) / b=1-a$, we have that

$$
a=1+\frac{p \log p}{1-p}
$$

On the other hand, due to the fixed point equation (8.7), the constant $P\left(T^{\prime}>\right.$ 1) $/ P(T>1)=\left(1-p^{\prime}\right) /(1-p)$ simplifies into $\alpha$.

Therefore, Theorem 8 is proved when we show that

$$
\alpha^{2}\left(1-2 a^{\prime}\right)>1-2 a
$$

Using the fixed point equation (8.7), we find that $a^{\prime}=1+p \log p^{\prime} /\left(\left(\alpha^{2}\right)(1-p)\right)$, so that the above reduces to

$$
-\alpha^{2}-2 \frac{p \log p^{\prime}}{1-p}>-1-2 \frac{p \log p}{1-p}
$$

or

$$
\frac{2 p \log \alpha}{1-p}>\alpha^{2}-1
$$

Using one more time the fixed point equation, one gets

$$
2 p \log \alpha>\alpha-1
$$

Since $\log (1-x)>-x$, it is therefore enough to show

$$
-2 p(1-\alpha)>\alpha-1 \text { or } 2 p<1
$$

which is precisely the condition that the branching process is supercritical.
8.6. Fragmentations in the supercritical range. In the previous section we have computed asymptotics for the expected number of clusters in the graph resulting from the superposition of the cycle structures of $\sigma$ and $\pi$. We now need to show that at the end of the walk $\left(\sigma_{r}, r=0, \ldots, a n\right)$, there are no more than $o(n)$ additional cycles that have been generated by fragmentation, compared to the number of clusters of $\Gamma$.

To do this, we use once again the dynamic point of view adopted to deal with the subcritical regime. Let $\tau_{1}, \ldots, \tau_{a n}$ be the decomposition of $\pi$ in product of an transpositions as evoked earlier, and let $\sigma_{r}=\sigma \cdot \tau_{1} \ldots \tau_{r}$.

Lemma 13. For each $1 \leq r \leq a n$, the expected number of cycles in $\sigma_{r}$ generated by fragmentation is $O\left(n^{1 / 2}\right)$.

This is similar to the Erdős-Renyi case of Berestycki and Durrett [3], Theorem 3. Lemma 13 does not claim that the number of fragmentations itself is $O\left(n^{1 / 2}\right)$, but that the number of extra cycles generated by fragmentation is $O\left(n^{1 / 2}\right)$. Just like in the Erdős-Renyi case, many of the cycles that are fragmented get reabsorbed by large components fairly quickly.

Proof of Lemma 13. There can never be more than $n^{1 / 2}$ cycles of size larger than $n^{1 / 2}$. On the other hand, by Lemma 11 , the probability that $\tau_{t}$ will create a fragment of size smaller than $n^{1 / 2}$ is at most $n^{1 / 2} / n=n^{-1 / 2}$. Therefore the expected number of such fragmentations is at most $a n \cdot n^{-1 / 2}=O\left(n^{1 / 2}\right)$.

At this point, Theorem 8 is proved.
Acknowledgments. I am very grateful to Laurent Saloff-Coste for introducing me to Gromov hyperbolic spaces and to Jim Pitman for pointing out the reference to Kolchin's representation theorem. I wish to thank Rick Durrett particularly for his help at several important stages of this work and for his support during my stay at Cornell University.

## REFERENCES

[1] Arratia, R., Barbour, A. D. and Tavaré, S. (2003). Logarithmic Combinatorial Structures: A Probabilistic Approach. European Mathematical Society Publishing House, Zürich. MR2032426
[2] Athreya, K. B. and Ney, A. (1972). Branching Processes. Springer, New York. MR0373040
[3] Berestycki, N. and Durrett, R. (2006). A phase transition in the random transposition random walk. Probab. Theory Related Fields. To appear.
[4] Dembo, A. and Zeitouni, O. (1993). Large Deviations Techniques and Applications. Jones and Bartlett, Boston, MA. MR1202429
[5] Durrett, R. (2004). Probability: Theory and Examples, 3rd ed. Duxbury Press, Belmont, CA. MR1609153
[6] Durrett, R. (2005). Random Graphs. In preparation.
[7] Gromov, M. (1987). Hyperbolic groups. Math. Sci. Res. Inst. Publ. 8 75-263. MR0919829
[8] Janson, S., Luczak, T. and Ruczinski, A. (2000). Random Graphs. Wiley, New York MR1782847
[9] Pitman, J. (2002). Poisson-Dirichlet and GEM invariant distributions for split-andmerge transformations of an interval partition. Combin. Probab. Comput. 11 501-514. MR1930355
[10] Pitman, J. (2006). Combinatorial Stochastic Processes. École d'Été de Probabilités de SaintFlour XXXII. Lecture Notes in Math. 1875. Springer, Berlin.
[11] Prat, J. J. (1971). Étude asymptotique du mouvement Brownien sur une variété Riemanienne à courbure négative. C. R. Acad. Sci. Sér A-B 272 A1586-A1589. MR0279889
[12] Schramm, O. (2005). Composition of random transpositions. Israel J. Math. 147 221-243. MR2166362

Department of Mathematics University of British Columbia 1984 Mathematics Road—RoOM 121
Vancouver, British Columbia
CANADA V6T 1Z2
E-MAIL: nberestycki@math.ubc.ca


[^0]:    Received December 2004; revised September 2005.
    ${ }^{1}$ Supported in part by Rick Durrett's joint NSF-NIGMS Grant DMS-02-01037.
    AMS 2000 subject classifications. Primary 60G50, 60K35, 60D05; secondary 60C05.
    Key words and phrases. Random walks, Gromov hyperbolic spaces, phase transition, random transpositions, random graphs, Cayley graphs.

