Discrete Mathematics and Theoretical Computer Science **AC**, 2003, 9–16

Random Infinite Permutations and the Cyclic Time Random Walk

Omer Angel

Dept. of Math., Weizmann Institute, Rehovot, Israel
omer@wisdom.weizmann.ac.il

The random stirring process is a natural random walk on the set of permutations of the vertex set of a graph. The cyclic time random walk is a self interacting random walk on a graph. It is influenced by its past, in that it is constrained to repeat its past choices if it returns to a previously visited edge after a multiple of some period of time. The two models are fundamentally equivalent to each other as well as to a certain coalescence and fragmentation process.

We show that when the graph is the d-regular tree there is a phase transition for these models. Infinite orbits appear in the random permutations, and the random walk becomes transient as the parameter (period) grows. Moreover, it also shown that such a transition occurs at $T_c = d^{-1} + O(d^{-2})$.

Keywords: Self interacting random walk, Random permutation, Phase transition

1 Introduction

The random stirring process (RSP) is a a process taking values in the group of permutations of the vertex set of some graph G. It was first introduced in [3], where it is related to the quantum Heisenberg ferromagnetic model. For any set V, a random permutation process σ_t for $t \geq 0$ taking values in the group S_V of permutations of V may be defined as follows: Begin with σ_0 which is the identity permutation. Repeatedly, a pair of elements of V is chosen randomly, and are exchanged. In the greatest generality there is no restriction on how the pairs are chosen. However, if the pairs are chosen independently of each other according to some fixed distribution, then the resulting random permutation process is seen to have i.i.d. increments (where the increment over an interval (s,t) is $\sigma_t\sigma_s^{-1}$). Thus this process is a caseof a (weighted) random walk on the Cayley graph of the permutation group S_V where neighbors differ by a single exchange.

Hereafter, a permutation of an infinite set is a bijection between the set and itself, and must be invertible, though it may have an infinite number of non-stationary points. An element $\sigma \in S_V$ partitions the underlying set into orbits, which are the finite cycles of σ (including fixed points) and infinite paths on which σ acts as a simple shift. E.g. the permutation on $\mathbb{Z}$ which fixes the odd integers and maps $2n$ to $2n+2$ has a finite orbit of size 1 for each odd integer and a single infinite orbit containing the even integers in their natural order.

To define the process rigorously a continuous time framework is used. Each pair of elements of V has some rate associated with it, and the pair is chosen according to a Poisson process with that rate. Even so when the set V is infinite, it is possible for the process to be ill-defined, and further constraints need to be placed on the rates in order to ensure that the resulting permutation at any time t is defined everywhere and is invertible.

We are interested in a special case of the above process where the set V is the vertex set of some graph G, so that some geometry is introduced in the process. We assume that the pairs are chosen from among the edges of G. In particular, our main result will concern the case where G is an infinite regular tree. Furthermore, we will assume that each edges is chosen with rate 1, i.e. each edge is associated with an independent Poisson process of density 1 on $\mathbb{R}^+$. If t is a point of the Poisson process associated with an edge e, then we say that e rings at time t.

To define σ_T formally, its values are traced as t moves from 0 to T. Fix some vertex v, and consider the evolution of $X_t = \sigma_t(v)$ as t changes. If at time t an edge e rings that is not incident on X_t then the ring has no effect on X_t. If an edge incident on X_t rings, then at the time of the ring the path jumps across the edge. It is clear that X_t is non other then the simple random walk on G in continuous time. For reasonable graphs (e.g. bounded degrees are sufficient), for any vertex v and any time t $\sigma_t(v)$ is indeed some vertex of G. Since $\sigma^{-1}(v)$ may be traced in the same manner, σ_t is invertible. We will not be concerned here with general sufficient conditions for X_t to be defined for all t.

The group S_V of permutations has a natural topology on it induced by convergence of values of the permutation and of its inverse (i.e. a sequence σ_n converges to σ if for every v and all $n > n_0(v)$, $\sigma_n(v) = \sigma(v)$ and $\sigma_n^{-1}(v) = \sigma^{-1}(v)$). Since a.s. each edge rings at a discrete set of times, for any t and any set of edges there is some ε so that the edges in the set do not ring in the interval $(t, t+\varepsilon]$. It follows that X_t is right-continuous ($X_t = X_{t^+}$) and so is the process σ_t. A Lévy process is a random function on $\mathbb{R}^+$ with i.i.d. increments and right continuous samples, so the random stirring is a special case of a *permutation valued Lévy process*.

The main question regarding the random stirring process concerns the cycle structure of the resulting permutation, and specifically the existence of infinite orbits of σ.

Question 1. *Let $Orb_\sigma(v)$ be the orbit in σ containing a vertex v. What can be said about the function $t \to \mathbb{P}(|Orb_{\sigma_t}(v)| = \infty)$, in the case of the RSP on a graph or for general permutation valued Lévy processes?*

For reasonable graphs (e.g. with bounded degrees), it is straightforward to see that $\sigma_t(v)$ is very unlikely to depend on the ringing times of edges far from v (as the random walk is unlikely to reach them by time t). It follows that the existence of infinite cycles may be approximated by tail events, and thus has probability 0 or 1 for any graph and any t. We define the critical T_c by $T_c = T_c(G) = \inf\{t \mid \sigma_t \text{ a.s. has infinite orbits}\}$.

A simple argument shows that $T_c > 0$ for many graphs: The ringing times of each edge are given by a Poisson process and so the probability that an edge does not ring by time t is e^{-t}, independently from all other edges. If the critical probability for bond percolation on G satisfies $p_c > 0$, and if t is small enough that $1 - e^{-t} < p_c$, then edges that ring at least once form only finite connected components. Since σ_t necessarily maps each such component to itself, for small t, σ_t will only have finite orbits. Thus we have proved

Theorem 2. $T_c(G) \geq -\ln(1 - p_c(G))$, *where $p_c(G)$ is the critical probability for bond percolation on G.*

However, this does not exclude the possibility that $T_c = \infty$ and σ_t has only finite orbits for any t. Our main result gives an upper bound on T_c in the case that G is a regular tree.

Theorem 3. *If G is a d-regular tree for $d \geq 5$, then there exists an interval $I_d = [a_d, b_d]$ such that σ_t a.s. has infinite orbits for $t \in I_d$. Moreover, for any $\varepsilon > 0$, for large enough d the above is satisfied with $I_d = [d^{-1} + (\frac{13}{6} + \varepsilon)d^{-2}, \ln 3 - \varepsilon]$.*

The reason we exclude $d = 3$ and $d = 4$ has little to do with the model itself, but is only a technical consequence of the given proof. A trivial lower bound on trees given by Theorem 2 is $\ln\left(\frac{d-1}{d-2}\right) = d^{-1} + 3/2d^{-2} + \ldots$ Thus the two bounds for T_c are very close to each other.

Corollary 4. *For d-regular trees*

$$\frac{3}{2} \leq \liminf_{d\to\infty} d^2(T_c - d^{-1}) \leq \limsup_{d\to\infty} d^2(T_c - d^{-1}) \leq \frac{13}{6}.$$

In the next section we describe the random walk in cyclic time (CTRW). In Section 3 the equivalence between the random stirring process and the CTRW is shown, as well as the connection to a coalescence and fragmentation model. The former equivalence was shown in [3], and is repeated here for completeness. These equivalences lead to some conjectures regarding the behavior of the process on various graphs. In Section 4 we give the proof of Theorem 3.

2 The Cyclic-Time Random Walk

The simple random walk in continuous time on a graph G is defined as by starting a particle (walker) at time 0 in some vertex v. For each edge incident on the walker's current position, the walker crosses that edge at rate 1. We assume as before that G is such that the particle a.s. makes only a finite number of steps in any finite time (this is the same weak condition that was needed in the definition of the stirring process). The SRW may equivalently be defined by having a Poisson process with density 1 on each edge, and letting the particle cross an edge when it rings. The cyclic time random walk (denoted CTRW for short, though it should not be mixed with the Continuous Time Random Walk) is a self interacting random walk where the interaction stems from a simple modification of the Poisson process of each edge. The ringing times of distinct edges are still independent from each other.

We now define the cyclic-time random walk with period T on a graph G. To each edge corresponds an independent Poisson process on $[0,T]$ which gives the edge's ringing times in that interval. The set of ringing times is then repeated periodically to get for each edge a set of times in $\mathbb{R}^+$. The above construction of the SRW in continuous time is then used with the periodic sets of ringing times. The only parameter in the definition is T, and the behavior of the CTRW may depend on T qualitatively.

A first simple observation is that a walker at vertex v at time $t + kT$ for an integer k will traverse exactly the same path as a walker at v at time t. Based on this, we will not always distinguish between times that differ by a multiple of T, and treat all times as being in the cycle $\mathbb{R}/T\mathbb{Z}$.

The CTRW as a random walk on the graph is not a Markov chain, since it is influenced by its past in complex manners, and not just through its present position. However, the CTRW is reversible in another sense: If we are given the ringing times and the particle's position at some time, it is possible to retrace the past path of the particle, i.e. the past of a particle is determined uniquely by its location at any given time. The path may be found by "reversing" the direction of time, and using the same procedure as for the regular CTRW. In particular, the probability of observing a path traversed during a certain time interval is equal to the probability of observing the reverse of the path.

An easy consequence of the definition of the CTRW is that if the walker returns to some previously visited vertex after a time that is a multiple of T, then it will necessarily traverse the same closed loop again and again, and in particular (since the walk is reversible) return infinitely often to the origin. In this case the path will visit only a finite set of vertices of G. Since until time T the CTRW is equivalent to the SRW, there is always a positive probability that the walker is at its starting vertex at time T and of this repetitious behavior occurring.

It is clear from the above and from the fact that a.s. each edge rings only at finitely many times modulo T, that the following are equivalent:
– The walker returns to the origin with probability 1.
– The walker a.s. returns infinitely often to the origin.
– A.s. only finitely many vertices are visited.
– A.s. the path has period kT for some $k \in \mathbb{N}$.
If these properties hold then we say that the CTRW with period T on G is *recurrent*, otherwise we say it is *transient*.

If the underlying graph G is recurrent, it seems reasonable that the CTRW will also be recurrent for any value of T. However, if G is transient the expected behavior is more interesting. Note that (as for the SRW) the dichotomy of transient and recurrent behavior does not depend on the starting vertex:

Proposition 5. *If the CTRW with period T on a connected graph G is transient when started at one vertex, then it is transient when started on any other vertex as well.*

The proof will be given in Section 3 using the connection with the random stirring process described there. Suppose on the other hand that the SRW on G is transient and that T is very large. Since up to time T the CTRW is identical to the simple RW, by time T the CTRW is likely to be far from the origin. However, as long as the CTRW is not at a vertex adjacent to its position at time $t-kT$ for some integer k, the CTRW will keep on making moves independent of its past. Thus the CTRW is likely to still be transient.

Conjecture 6. *The CTRW on G is recurrent for $T < T_c$ and transient for $T > T_c$. If G is recurrent then $T_c = \infty$, i.e. the CTRW on is recurrent for all T.*

If the SRW is identified with the case $T = \infty$ then this conjecture states that the transience property is non-decreasing for T in $[0,\infty]$. Note that the converse does not hold, since there exist transient graphs with $p_c = 1$ (e.g. in a binary tree replace edges at height n by n-long paths). Whenever $p_c = 1$, Theorem 2 implies $T_c = \infty$.

3 Equivalence of the Models

As the above discussion hints, there is a fundamental connection between the CTRW and the random stirring process. Note that since both are defined as deterministic functions of collection of Poisson processes, there is a natural coupling between the CTRW with period T and σ_T, where they are both constructed using the same Poisson processes.

Proposition 7. *In the natural coupling between the CTRW with period T and σ_T, the CTRW started at v returns to v infinitely often iff $Orb_{\sigma_T}(v)$ is finite.*

Proof. Using the same Poisson processes, we can trace the path of a walker starting at any vertex of G. If a walker starts at v, then at time T it will be at $\sigma_T(v)$. If the pattern of ringing is repeated with period

T, then a walker in v at time kT will be at $\sigma_T(v)$ at time $(k+1)T$. Thus a walker starting at v will reach $\sigma_T^k(v)$ at time kT.

If $|\mathrm{Orb}_{\sigma_T}(v)| = k < \infty$ then at time kT the walker will return to v. On the other hand, if the orbit is infinite, then at each multiple of T the walker will be at a new vertex. □

Corollary 8. *Let G be connected, then the CTRW on G is transient iff σ_T has infinite orbits with positive probability.*

Proof of Proposition 5. Since G is connected it is sufficient to prove the Proposition for neighboring vertices u, v. Let A be the event that v has an infinite orbit in σ_T, and assume $\mathbb{P}(A) = \delta > 0$. Let B_ε be the event that no edge incident on either u or v rings after time $T - \varepsilon$. Fix $\varepsilon > 0$ small enough that $1 - \mathbb{P}(B_\varepsilon) < \delta/2$, and then $\mathbb{P}(A \cap B_\varepsilon) > \delta/2$. Assume that the orbit of u is a.s. finite.

Let C be the event that the edge (u, v) rings exactly once after time $T - \varepsilon$. By the nature of Poisson processes, conditioning on C is equivalent to replacing the ringing times of (u, v) after time $T - \varepsilon$ (if any) by a single ring at a uniform time in $[T - \varepsilon, T]$. However, on the event B_ε an exchange of u and v commutes with any of the exchanges that contribute to σ_T after time $T - \varepsilon$, and hence the resulting σ'_T is simply σ_T followed by an exchange of u and v. It follows that conditioned on C, with probability at least $\delta/2$ the orbit of u is infinite, contradicting our assumption. □

Using Proposition 7 to translate Conjecture 6 to the framework of random permutations we get the following conjecture:

Conjecture 9. *If G is recurrent then for all T, σ_T a.s. has only finite orbits. If G is transient then there is some T_c such that a.s. σ_T has only finite cycles for $T < T_c$ and a.s. has some infinite orbits for $T > T_c$.*

3.1 Coalescence and Fragmentation

It is interesting to note one more incarnation of the above model (which is closely related to the random stirring process). The connection arises from observing how the orbits of σ_t evolve with t. The orbits are effected when edges ring, and each edge rings at rate 1.

When an edge (u, v) rings with u and v on distinct orbits, the result is that the two orbits connect at u and v and coalesce to a single larger orbit. When an edge connecting two vertices within the same orbit rings the orbit breaks down at the two vertices and is separated into two orbits.

Initially, σ_0 is the identity and each orbit consists of a single point. At any time each orbit is not just a set of vertices, but is endowed with a cyclic order where v is followed by $\sigma_t(v)$. If the orbit is infinite then the order is equivalent to $\mathbb{Z}$. If an edge (u, v) rings at time t, with u and v in distinct orbits of σ_{t^-}, the two orbits coalesce. The action of σ_t on the new orbit is identical for all vertices, except that $\sigma_t(\sigma_{t^-}(u)) = v$ and vice versa. If one of the two orbits of σ_{t^-} is infinite, then the new orbit is also infinite. An exception occurs when both orbits of σ_{t^-} are infinite. In this case they do not simply coalesce, and after the exchange the result is two infinite but different orbits.

As long as the orbits are finite, if the cyclic order within orbits is ignored, then the coalescence element is equivalent to the model where clusters coalesce at a rate equal to the number of edges between them. This is exactly the process describing coalescence of percolation clusters when each edge is initially closed and opens at rate 1.

In order to understand the fragmentation element the order within orbits can not be ignored. An orbit breaks into two orbits whenever an edge connecting two vertices within the orbit rings. The rate of fragmentation of an orbit equals the number of edges spanned by the vertices of the orbit. The cyclic structure determines which orbit each vertex belongs to after the break up.

Thus the CTRW is equivalent to the above coalescence/fragmentation model where there is a geometric aspect to both parts of the process, and some geometric information stored within each cluster. The question of transience and recurrence of the CTRW becomes in this formulation the question of appearance of infinite clusters at some finite time when starting with the state where each vertex is a cluster of its own.

4 The CTRW on Trees

In this section we present the proof of Theorem 3 regarding the stirring process and CTRW on regular trees. The simple tree geometry allows for a detailed analysis of the transience of the CTRW. The key idea of the proof is to find some pattern of ringing times for edges in some sub-tree that forces the particle never to return to its starting vertex. The pattern is defined by simple local constraints on the ringing times of edges near each vertex of the sub-tree. Consequently, the sub-tree satisfying the constraints will be shown to be a Galton–Watson tree. If the Galton–Watson tree is infinite, then it is the sub-tree we need. Unfortunately, the pattern we define only works for bounded values of T, and not for large T. This is a bit surprising since on the face of it it would appear that a larger period would only help transience.

Proof of Theorem 3. Call the starting vertex the root of the tree. Every other vertex has a father which is its nearest neighbor in the direction of the root. A vertex v is said to be *good* if the edge connecting v and its father rings exactly once. Thus each good vertex has a unique time t_v in $[0,T]$ at which it may be reached from its father. A particle entering a good vertex from its father will either never return to the father, or return at the same time it arrived (modulo T). For the root we set $t_{\mathrm{root}} = 0$.

A good vertex v with a good father u is said to be *covered* by a sibling v' (of v) if the ringing times of the edge (u,v') separate t_u and t_v cyclically modulo T. This means that during the interval $[t_u, t_u+T]$ the edge (u,v') rings both before and after the single time that (u,v) rings. The idea is that a particle arriving at u from u's father might not visit v and return to u's father by passing to v' at the first time (u,v') rings and returning from v' after (u,v) has already rang. A good vertex v with a good father u is uncovered if it is not covered by any of its siblings.

We claim that if v is uncovered and the particle reaches u at all, then one of two cases must hold. Either the particle will also visit v at some future time, or the particle will pass at some time from u to another vertex and never return to u again. In the latter case the particle's path is of course necessarily transient (this strongly utilizes the tree geometry).

To prove this claim, note that the particle first reaches u at t_u (modulo T). Suppose the second case fails to occur, so that whenever the particle leaves u it is guaranteed to return back to u. Since v is uncovered, any edge incident on u that rings during the cyclic interval $I = [t_u, t_v]$ rings only during that interval. It follows that as long as the particle does not pass to v it can only be at u at times inside that interval. Indeed, if the particle leaves u along some edge (u,v'') that rings during I, then since G is a tree it can only return along the same edge, and so return inside the interval I. Since a.s. all edges ring a finite number of times, and since the particle cannot enter some finite cycle that does not involve returning to u at t_u, eventually the particle will necessarily be in u at time t_v and therefore enter v.

Let N_u be the number of good uncovered offspring of a good vertex u. Since the definition of a good uncovered vertex is invariant to cyclically shifting the ringing times of all edges at u, the distribution of N_u does not depend on t_u. It follows that N_u is independent of N_v for all other good vertices. Thus the good uncovered connected component of the root is a Galton–Watson tree (with the exception that the root has d potential offspring rather than $d-1$, and so has a different offspring distribution). If the expected number of offspring is greater than 1, then with positive probability there is an infinite path of good uncovered vertices from the root (see [1] for results on Galton–Watson trees), and by the above claim the particle's path must be transient.

It remains to estimate the expectation of N_u for a good vertex u. Assume w.log. $t_u = 0$. The expectation is $d-1$ times the probability that a specific child v of u is good and uncovered. The edge (u,v) rings exactly once with probability e^{-T}, and conditioned on that, its ringing time is uniform. If $t_v = a$ then the probability that v is covered by a sibling v' is $(1-e^{-a})(1-e^{a-T})$ and since each v has $d-2$ siblings, the expected number of uncovered edges above any good edge is

$$\begin{aligned}\mathbb{E}N_u &= (d-1)e^{-T}\int_0^T \left(1-(1-e^{-a})(1-e^{a-T})\right)^{d-2} da \\ &= (d-1)e^{-T}\int_0^T \left(e^{-a}+e^{a-T}-e^{-T}\right)^{d-2} da.\end{aligned} \tag{1}$$

For any fixed d it is straightforward to calculate this integral. For $5 \le d \le 8$ we find numerically that Theorem 3 holds with $I_5 = [0.40, 1.11]$, $I_6 = [0.28, 1.23]$, $I_7 = [0.22, 1.28]$, $I_8 = [0.18, 1.30]$. For larger values of d the following gives a valid interval I_d. Note

$$e^{-a}+e^{a-T}-e^{-T} \ge \max(e^{-a}, e^{a-T})$$

for $a \in [0,T]$. Using this and symmetry of the integrand gives

$$\begin{aligned}\mathbb{E}N_u &\ge 2(d-1)e^{-T}\int_0^{T/2} e^{-a(d-2)} da = \frac{2d-2}{d-2}(e^{-T}-e^{-Td/2}) \\ &\ge 2(e^{-T}-e^{-Td/2}).\end{aligned}$$

It is easy to check (using monotonicity of the bound) that the last bound is greater that 1 for any $d \ge 9$ and any $T \in [4/d, 1/2]$. The asymptotics of $\mathbb{E}N_u$ as $d \to \infty$ are $\mathbb{E}N_u = \frac{2}{e^T-1} + O(d^{-1})$ and so for large d the interval may be extended up to $\ln 3 - \varepsilon$.

It remains to be shown that this interval can be extended to include still smaller values of T. For large d and $T < 4/d$ we use the expansion $1 - x + \frac{x^2}{2} - \frac{x^3}{6} < e^{-x} < 1 - x + \frac{x^2}{2}$ for all $x > 0$. Plugging this into (1) gives for $T < 4/d$ and some c

$$\begin{aligned}\mathbb{E}N_u &> (d-1)e^{-T}\int_0^T \left(1 - a(T-a) - \frac{T^3}{6}\right)^{d-2} da \\ &> (d-1)e^{-T}\int_0^T \left(1 - da(T-a) - cT^2\right) da \\ &= (d-1)e^{-T}(T - dT^3/6 - cT^3/3).\end{aligned}$$

For $T = d^{-1} + \gamma d^{-2}$ this reduces to $1 + (\gamma - 13/6)d^{-1} + O(d^{-2})$. This implies that for any $\gamma > 13/6$, large enough d, and $T \in [d^{-1} + \gamma d^{-2}, 4d^{-1}]$, we have $\mathbb{E}N_u > 1$. $\square$

5 Open Problems

We would like to finish by pointing out a few open problems concerning the introduced class of models.

Problem 10. *Corollary 4 gives the first term in the asymptotic behavior of T_c on regular trees, and bounds the error term. What is the correct asymptotics for T_c?*

It would appear that neither the upper nor the lower bound is correct. The largest error term in both bounds arise from edges that ring exactly twice. It is possible to incorporate such edges into either proof to tighten the bounds. Edges that ring twice in quick succession are likely not to make any change for the CTRW, but removing such edges may shift the Galton–Watson process from being super-critical to being sub-critical. Similarly, the upper bound can be tightened. It is likely that with a precise enough understanding of the effect of edges that ring twice it is possible to reduce the gap to $O(d^{-3})$.

Similarly, using a somewhat richer structure that utilizes also edges that ring twice it should be possible to extend the result also to $d = 3,4$.

Problem 11. *What happens on $\mathbb{Z}^d$ or other graphs?*

The case $d = 1$ is trivial. Conjecture 6 would imply there is a significantly different behavior for $d = 2$ and $d > 2$. Is it possible to translate the proof from trees to large dimensional lattices? Can It be translated to hyperbolic lattices?

When $d > 4$ even a SRW is likely not to intersect its past. Can this be used to understand the CTRW, as it was used to understand the self avoiding walk [2]?

Problem 12. *What can be said about the cycle structure of general permutation Lévy processes, and in particular the appearance of infinite orbits?*

For a process σ_t, let p_t be the probability that σ_t has infinite orbits. p_t is not necessarily 0 or 1. For example acting on $\mathbb{Z}$, set $\sigma_t(n) = n + X_t$ where X_t is a SRW on $\mathbb{Z}$ in continuous time. It is clear that in this case $0 < p_t < 1$ for all $t > 0$.

What may be said about the possible functions $t \to p_t$? Is p_t necessarily monotone? Is it monotone for the RSP on a graph? Is it true that if $p_t = 0$ for all $t < \varepsilon$ then there is some t_c such that $p_t = 0$ for $t < t_c$ and is positive for $t > t_c$? Is it true that $p_t = 1$ for $t > t_c$?

Acknowledgements

The author would like to thank Itai Benjamini for pointing out the CTRW model to him and Yuval Peres for helpful comments.

References

[1] R. Lyons and Y. Peres, Probability on Trees and Networks. `http://mypage.iu.edu/~rdlyons/prbtree/prbtree.html`.

[2] T. Hara and G. Slade, Self-avoiding walk in five or more dimensions. *Comm. Math. Phys.* **147** (1992), no. 1, 101–136.

[3] B. Toth, Improved lower bound on the thermodynamic pressure of the spin $1/2$ Heisenberg ferromagnet. *Lett. Math. Phys.* **28** (1993) no. 1, 75–84.