

On the cover time of dense graphs

Colin Cooper*

Alan Frieze[†]

Wesley Pegden[‡]

June 3, 2020

Abstract

We consider arbitrary graphs G with n vertices and minimum degree at least δn where $\delta > 0$ is constant.

(a) If the conductance of G is sufficiently large then we obtain an asymptotic expression for the cover time C_G of G as the solution to an explicit transcendental equation.

(b) If the conductance is not large enough to apply (a), but the mixing time of a random walk on G is of a lesser magnitude than the cover time, then we can obtain an asymptotic deterministic estimate via a decomposition into a bounded number of dense subgraphs with high conductance.

(c) If G fits neither (a) nor (b) then we give a deterministic asymptotic $(2+o(1))$ -approximation of C_G .

1 Introduction

Let $G = (V, E)$ be a connected graph with vertex set $V = [n] = \{1, 2, \dots, n\}$ and an edge set E of m edges. In a simple random walk W on a graph G , at each step, a particle moves from its current vertex to a randomly chosen neighbour. For $v \in V$, let C_v be the expected time taken for a simple random walk starting at v to visit every vertex of G . The *vertex cover time* C_G of G is defined as $C_G = \max_{v \in V} C_v$. The (vertex) cover time of connected graphs has been extensively studied. It is a classic result of Aleliunas, Karp, Lipton, Lovász and Rackoff [2] that $C_G \leq 2m(n-1)$. It was shown by Feige [14], [15], that for any connected graph G , the cover time satisfies $(1 - o(1))n \log n \leq C_G \leq (1 + o(1))\frac{4}{27}n^3$. As an example of a graph

*Research supported in part by EPSRC grant EP/M005038/1

[†]Research supported in part by NSF grant DMS1362785

[‡]Research supported in part by NSF grant DMS1363136

achieving the lower bound, the complete graph K_n has cover time determined by the Coupon Collector problem. The *lollipop* graph consisting of a path of length $n/3$ joined to a clique of size $2n/3$ gives the asymptotic upper bound for the cover time.

It follows from [2] that there is a very simple randomised algorithm for estimating the cover time. Simply execute enough random walks so that the average cover time can be used as an estimate. It is much more challenging to estimate the cover time deterministically in polynomial time. A theorem of Matthews [20] gives a deterministic $O(\log n)$ approximation. This was improved to $O((\log \log n)^2)$ by Kahn, Kim, Lovász and Vu [16]. In a breakthrough, Ding, Lee and Peres [13] improved this to $O(1)$ using a remarkable connection between the cover time and Gaussian free fields. Subsequently, Ding [12] has improved the factor of approximation to $1 + o(1)$, as $n \rightarrow \infty$ for trees and bounded degree graphs. Zhai [24] has recently shown that if the maximum hitting time is asymptotically smaller than the cover time then the approximation ratio is $1 + o(1)$, implying the results of [13] and [12]. An important point to note here is that Meka [21] gives a polynomial time approximation scheme for finding the supremum of a Gaussian process. This is what provides the computational underpinning for the results of [13], [12] and [24]. We note that none of these results give an explicit value of the cover time as a function of the number of vertices n or imply a deterministic polynomial time approximation scheme for the cover time.

The first two authors of this paper have studied the cover time of various models of a random graph, see [5, 6, 8, 9]. The main tool in their analysis has been the “First Visit Lemma”, see Lemma 4. In this paper we see how this lemma can be used deterministically to give good estimates of the cover time of dense graphs when the mixing time is asymptotically smaller than the cover time.

Let $\mathcal{G}(n, \theta)$ denote the set of connected graphs with vertex set $[n]$ and minimum degree at least θn . Our first result deals with the simplest case, where the mixing time of a random walk on our graph is sufficiently small. Subsequent theorems will consider more general cases.

Notation: The degree sequence of the graph $G = (V, E)$, $|V| = n$, will always be $\mathbf{d} = (d_1, d_2, \dots, d_n)$ so that $2m = \sum_{i=1}^n d_i$. For $S \subseteq V$ we let $d(S) = \sum_{i \in S} d_i$ and $e(S) = |\{v, w\} \in \binom{S}{2} \cap E|$.

For two sequences A_n, B_n we write $A_n = (1 \pm \epsilon)B_n$ if $(1 - \epsilon)B_n \leq A_n \leq (1 + \epsilon)B_n$ for n sufficiently large. For two sequences A_n, B_n we write $A_n \approx B_n$ if $A_n = (1 + o(1))B_n$ as $n \rightarrow \infty$. We will write $A_n \gg B_n$ or $B_n \ll A_n$ to mean that $A_n/B_n \rightarrow \infty$ as $n \rightarrow \infty$.

For $S, T \subseteq V$, $S \cap T = \emptyset$ let $e(S, T) = e_G(S, T) = |\{v, w\} \in E : v \in S, w \in T|$, let $\bar{S} = V \setminus S$ and

$$\Phi(S) = \Phi_G(S) = \frac{e(S, \bar{S})d(V)}{d(S)d(\bar{S})}. \quad (1)$$

The *conductance* $\Phi(G)$ of G is given by

$$\Phi(G) = \min_{0 < d(S) \leq m} \Phi_G(S).$$

We will make our walk *lazy* and ergodic by adding a loop at each vertex so that the walk stays put with probability $1/2$ at each step. This has the effect of (asymptotically) doubling the cover time and the extra factor of two can be discarded. (Ergodicity only requires a small probability of staying in place, but laziness allows us to use conductance to estimate the mixing time. See (18).) A simple random walk has a steady state of

$$\pi_i = \frac{d_i}{2m}, i \in [n] \text{ and if } G \in \mathcal{G}(n, \theta) \text{ then } \frac{\theta}{n} \leq \pi_i \leq \frac{1}{\theta n}. \quad (2)$$

Next let

$$F(t) = \sum_{v \in V} \frac{e^{-\pi_v t}}{\pi_v} \quad \text{and so} \quad F'(t) = - \sum_{v \in V} e^{-\pi_v t}. \quad (3)$$

Note that F is monotone decreasing and F' is monotone increasing and that F is convex. Next let

$$\psi = \frac{1}{\log^{2/3} n}. \quad (4)$$

Theorem 1. *Let $\epsilon > 0$ be arbitrary and suppose that $G \in \mathcal{G}(n, \theta)$ where $\theta = \Omega(1)$. Suppose that $\Phi = \Phi(G) \geq n^{-\theta\psi}$. Then there exists n_ϵ such that if $n \geq n_\epsilon$ then*

$$C_G = (1 \pm \epsilon)t^* \quad (5)$$

where t^* is the unique solution to $F'(t) = -1$, (see (3)).

Thus, if G is regular and satisfies the conditions of Theorem 1 then $C_G \approx n \log n$. Also, if G is regular of degree θn where $\theta > 1/2$, then the conditions of Theorem 1 will be satisfied. Indeed, the condition that $d(S) \leq m$ in the definition of conductance is equivalent to $|S| \leq n/2$ and then

$$\Phi(S) \geq \frac{(\theta - \frac{1}{2}) n |S| d(V)}{n |S| d(V)} = \theta - \frac{1}{2}.$$

Remark 1. *Note also, that while it may be difficult to compute $\Phi(G)$ exactly in deterministic polynomial time, we can approximate it to within an $O(\log n)$ factor using the algorithm of Leighton and Rao [19]. Thus if $\Phi(G) \gg n^{-\theta\psi} \log n$ then there is a deterministic polynomial time algorithm that verifies that G satisfies the conditions of Theorem 1 and gives a $(1 + \epsilon)$ -approximation to the cover time.*

The proof of Theorem 1 is given in Section 3 and closely follows the lines of the proofs for random instances.

Suppose that we start our walk $\mathcal{W}_u = (u = \mathcal{W}_u(0), \mathcal{W}_u(1), \dots, \mathcal{W}_u(t), \dots)$ at vertex u and that $P_u^{(t)}(x) = \mathbf{Pr}(\mathcal{W}_u(t) = x)$. Let

$$d(t) = \max_{u, x \in V} |P_u^{(t)}(x) - \pi_x|,$$

and let $T_{mix} = T_{mix}(\omega)$ be such that, for $t \geq T_{mix}$

$$\max_{u, x \in V} \left| \frac{P_u^{(t)}(x) - \pi_x}{\pi_x} \right| \leq \frac{1}{\omega} \quad (6)$$

where $\omega = \omega(n) \rightarrow \infty$. We will assume that

$$\omega = n^{3\theta\psi}. \quad (7)$$

If the conditions of Theorem 1 fail, then we partition the vertex set V into $O(1)$ subsets which satisfy the conditions of Theorem 1. If furthermore, our *mixing time*

$$T_{mix} = o(C_G)$$

then we will obtain a $(1 + \epsilon)$ -approximation to the cover time.

Remark 2. Note that by examining the powers of the transition matrix P , we can determine the mixing time $T_{mix}(u, \omega)$, $u \in V$ in deterministic polynomial time. We note that $T_{mix}(u) = O(n^3)$, as long as the accuracy needed in (6) is at most $1/\omega = e^{-poly(n)}$, see [18] (Proposition 10.28). In which case we only need to compute a $poly(n)$ power of P .

Theorem 2. Let $\epsilon > 0$ be arbitrary and suppose that $G \in \mathcal{G}(n, \theta)$ where $\theta = \Omega(1)$. Then in deterministic polynomial time we can find a partition of V into subsets $V_i, i = 1, 2, \dots, s$ where $s = O(1)$, where the induced subgraphs $G[V_i]$ satisfy the conditions of Theorem 1, and have cover time C_i which can be computed via Theorem 1.

Suppose furthermore that $T_{mix} = o(C)$, where C is given by

$$C = \max \left\{ \frac{C_i}{\pi(V_i)} : i \in [s] \right\}, \quad (8)$$

then $C_G = (1 \pm \epsilon)C$.

The construction of this partition is described in Section 4.

Finally, if T_{mix} is too large for Theorem 2 to apply then we do not have a nice expression for C_G , but instead we have

Theorem 3. Let $\epsilon > 0$ be arbitrary and suppose that $G \in \mathcal{G}(n, \theta)$ where $\theta = \Omega(1)$. Then in deterministic polynomial time we can compute an estimate $\overline{C}_G$ such that if $n \geq n_\epsilon$ then

$$\overline{C}_G \leq C_G \leq (2 + o(1))\overline{C}_G. \quad (9)$$

The proof of Theorem 2 uses a concentration inequality of Paulin [22], which requires a sufficiently small mixing time. The proof of Theorem 3 uses the partition of Theorem 2. It then describes how to use the transition matrix of the walk to give upper and lower estimates for the time needed to visit each V_i .

2 First Visit Lemma

Our main tool will be Lemma 4 below. The lemma has been used several times in the context of random graphs, see for example [5, 6, 8, 9]. We sharpen the proof to make it fit the current situation. Let G denote a fixed connected graph, and let u be some arbitrary vertex from which a walk $\mathcal{W}_u$ is started. Let $\mathcal{W}_u(t)$ be the vertex reached at step t , let P be the matrix of transition probabilities of the walk, and let

$$h_t = P_u^{(t)}(v) = \mathbf{Pr}(\mathcal{W}_u(t) = v).$$

It follows from e.g. Aldous and Fill [1], Lemma 2.20, that $d(t)$ satisfies $d(s+t) \leq 2d(s)d(t)$ which implies that

$$\max_{u,x \in V} |P_u^{(kT)}(x) - \pi_x| \leq 2^{k-1} (\max_{u,x \in V} |P_u^{(T)}(x) - \pi_x|)^k \leq 2^{k-1} \left(\frac{\pi_x}{\omega} \right)^k.$$

And because $d(t)$ is monotone decreasing in t , for $t \geq T = T_{mix}$ and $k = \lfloor t/T \rfloor$, we have

$$\max_{u,x \in V} \left| \frac{P_u^{(t)}(x) - \pi_x}{\pi_x} \right| \leq \frac{2^{k-1}}{\omega^k}. \quad (10)$$

Fix two vertices u, v . Let

$$H(z) = \sum_{t=T}^{\infty} h_t z^t \quad (11)$$

generate h_t for $t \geq T$.

Next, considering the walk $\mathcal{W}_v$, starting at v , let $r_t = \mathbf{Pr}(\mathcal{W}_v(t) = v)$ be the probability that this walk returns to v at step $t = 0, 1, \dots$. Let

$$R(z) = \sum_{t=0}^{\infty} r_t z^t$$

generate r_t . Our definition of return includes $r_0 = 1$.

For $t \geq T$ let $f_t = f_t(u \rightarrow v)$ be the probability that the first visit of the walk $\mathcal{W}_u$ to v in the period $[T, T+1, \dots]$ occurs at step t . Let

$$F(z) = \sum_{t=T}^{\infty} f_t z^t$$

generate f_t . Then we have

$$H(z) = F(z)R(z). \quad (12)$$

Finally, let

$$R_T(z) = \sum_{j=0}^{T-1} r_j z^j \text{ and } H_T(z) = \sum_{j=0}^{T-1} h_j z^j. \quad (13)$$

Now fix $u \neq v \in V$. For a large constant $K > 0$, let

$$\lambda = \frac{1}{KT}. \quad (14)$$

For $t \geq 0$, let $\mathcal{A}_t(v)$ be the event that $\mathcal{W}_u$ does not visit v in steps $T, T+1, \dots, t$. The vertex u will have to be implicit in this definition.

Lemma 4. *Suppose that*

(a) *For some constant $c > 0$, we have*

$$\min_{|z| \leq 1+\lambda} |R_T(z)| \geq c. \quad (15)$$

(b)

$$T\pi_v \leq \omega^{-1} = o(1). \quad (16)$$

Let $R_v = R_T(1)$. Then we can write

$$p_v = \frac{\pi_v}{R_v(1 + \xi_{v,1})} \text{ where } |\xi_{v,1}| = O(\omega^{-1}).$$

And then for all $t \geq T$,

$$\Pr(\mathcal{A}_t(v)) = \frac{1 + \xi_{v,2}}{(1 + p_v)^t} + o(Te^{-\lambda t/2}) \text{ where } |\xi_{v,2}| = O(\omega^{-1}). \quad (17)$$

Proof. The proof is very similar to that given in previous papers. We will defer its proof to an appendix. $\square$

Remark 3. *We will not have to verify (15) to use the theorem. It was shown in [11] that (15) follows from $R_v = O(1)$ and in our applications, $R_v = 1 + o(1)$.*

3 Proof of Theorem 1

Because our results require $n \geq n_\epsilon$ we can state inequalities in asymptotic terms. I.e. if we want to show that some parameters A_n, B_n satisfy $A_n \leq (1 + \epsilon)B_n$ then we can write $A_n \leq (1 + o(1))B_n$. Then if n is large enough, so that the $o(1)$ term is at most ϵ , then we are dealing with a bounded size problem, which can in principle, be dealt with by an exponential time algorithm.

We continue by computing parameters for use in Lemma 4. We begin with the mixing time $T = T_{mix}$. We use the following Cheeger inequality, see for example Levin, Peres and Wilmer [18], (13.6).

$$\max_{u, x \in V} \left| \frac{P_u^{(t)}(x) - \pi_x}{\pi_x} \right| \leq \frac{e^{-\Phi^2 t/8}}{\min_u \pi_u} \leq \theta^{-1} n e^{-\Phi^2 t/8}, \quad (18)$$

where the last inequality follows from (2). (We have $e^{-\Phi^2 t/8}$ instead of $e^{-\Phi^2 t/2}$ because our definition of Φ is larger than that defined in (7.8) of [18], but larger by a factor of at most 2.)

We can satisfy (6) if we take

$$T = \frac{8 \log(\omega n / \theta)}{\Phi^2} = \frac{8 \log(n^{1+3\theta\psi} / \theta)}{\Phi^2} \leq \frac{10\theta \log n}{n^{-2\theta\psi}} \leq n^{3\theta\psi}. \quad (19)$$

With this value of T we find that

$$T\pi_v \leq \frac{T}{\theta n} \leq \frac{n^{3\theta\psi}}{\theta n} \leq \omega^{-1} = n^{-3\theta\psi}. \quad (20)$$

We therefore find that (16) is satisfied.

Lemma 5. *Let t^* be as in Theorem 1. Then,*

$$(a) \quad F(t^*) = o(t^*). \quad (21)$$

$$(b) \quad n \log n \leq t^* \leq \theta^{-1} n \log n. \quad (22)$$

Proof. Now, by convexity,

$$\frac{1}{n} = \sum_{v \in V} \frac{1}{n} e^{-t^* \pi_v} \geq e^{-\frac{t^*}{n} \sum \pi_v} = e^{-t^*/n}.$$

and this implies the lower bound in (22).

Next observe that (2) implies

$$1 = \sum_{v \in V} e^{-t^* \pi_v} \leq n e^{-\theta t^* / n}$$

and this implies the upper bound in (22). Also we have that

$$\frac{F(t^*)}{t^*} \leq \sum_{v \in V} \frac{e^{-t^* \pi_v}}{\theta \log n} = \frac{1}{\theta \log n} = o(1),$$

as claimed in (21). □

3.1 Upper bound on C_G

We consider the walk $\mathcal{W}_u$ and write $T = T_{mix}(u)$ for the mixing time. We observe first that

$$1 \leq R_v \leq 1 + \frac{T}{\theta n} = 1 + o(1). \quad (23)$$

The inequality follows from the fact that if the walk $\mathcal{W}_v$ is not at v then the probability it moves to v at the next step is at most $1/\theta n$. The final claim can be seen from (20). Here we have ignored the self loops added to each vertex to make the chain lazy. As already mentioned, the addition of these loops multiplies the covertime by $(2 + o(1))$ and so to get the covertime we would multiply and then divide by this factor. Here we just acknowledge that it multiplies the mixing time by a factor $(2 + o(1))$, which can also be ignored in the above equation.

Let $T_{cov}(u)$ be the time taken to visit every vertex of G by the random walk $\mathcal{W}_u$. Let U_t be the number of vertices of G which have not been visited by $\mathcal{W}_u$ at step t . We note the following:

$$\Pr(T_{cov}(u) > t) = \Pr(U_t \geq 1) \leq \min\{1, \mathbf{E}(U_t)\}, \quad (24)$$

$$C_u = \mathbf{E}(T_{cov}(u)) = \sum_{t \geq 0} \Pr(T_{cov}(u) \geq t) \quad (25)$$

It follows from (24), (25) that for all t

$$C_u \leq t + 1 + \sum_{s > t} \mathbf{E}(U_s) \leq t + 1 + \sum_{v \in V} \sum_{s > t} \Pr(\mathcal{A}_s(v)). \quad (26)$$

Putting $t = t^*$, defined in (5), we see from (17) that

$$\begin{aligned}
C_u &\leq t^* + 1 + \sum_{v \in V} \sum_{s > t^*} \left(\frac{(1 + \xi_{v,2})}{(1 + p_v)^s} + o(Te^{-\lambda s/2}) \right) \\
&= t^* + 1 + \sum_{v \in V} \left(\frac{(1 + \xi_{v,2})}{p_v(1 + p_v)^{t^*+1}} + o(T^2e^{-\lambda t^*/2}) \right) \\
&= t^* + 1 + \sum_{v \in V} \left(\frac{(1 + \xi_{v,2}) \exp \{-(t^* + 1) \log(1 + p_v)\}}{p_v} + o(T^2e^{-\lambda t^*/2}) \right) \\
&= t^* + 1 + (1 + O(\omega^{-1})) \sum_{v \in V} \left(\frac{e^{-p_v t^* + O(p_v^2 t^*)}}{p_v} + o(T^2e^{-\lambda t^*/2}) \right). \tag{27}
\end{aligned}$$

Remark 4. Observe that the term $o(T^2e^{-\lambda t^*/2})$ is negligible, since $t^* = \Theta(n \log n)$ and $\lambda = \Omega(n^{-3\epsilon\theta})$. It is in fact at most $e^{-n^{1-3\epsilon\theta}}$ and we will assume always that ϵ is sufficiently small.

Now, because $t^* = \Theta(n \log n)$, we have, using (20) and (2),

$$p_v t^* = \left(1 + O\left(\frac{T}{\theta n}\right) + O\left(\frac{1}{n}\right) \right) \pi_v t^* = \pi_v t^* + O(n^{-1+3\theta\psi} \log n) \text{ and } p_v^2 t^* = O\left(\frac{\log n}{n}\right).$$

And so we can replace (27) by

$$C_G \leq t^* + 1 + (1 + O(\omega^{-1}))F(t^*) + O(e^{-n^{1-3\epsilon\theta}}) \leq \left(1 + \frac{2}{\theta \log n} \right) t^* \leq (1 + \epsilon)t^*, \tag{28}$$

after using (21).

Remark 5. Using Remark 4 we obtain a simpler upper bound:

$$\Pr(T_{cov}(u) \geq t) \leq \sum_{v \in V} \Pr(\mathcal{A}_t(v)) \leq (1 + o(1)) \sum_{v \in V} \frac{e^{-(1+o(1))t\pi_v}}{\pi_v}.$$

Putting $t = Kt^*$, we see that for any constant $L > 0$ there exists $K = K(L)$ such that

$$\Pr(T_{cov}(u) \geq Kt^*) \leq n^{-L}.$$

3.2 Lower bound

Now let T, u be as in Section 3.1.

$$t_1 = t^*(1 - \epsilon_1) \text{ where } \epsilon_1 = \frac{1}{\log^{1/2} n}. \tag{29}$$

Then let U_1 denote the set of vertices that have not been visited by $\mathcal{W}_u$ by time t_1 , and let $\mathcal{T} = \{\mathcal{W}_u(i) : 1 \leq i \leq T\}$. Then we have that

$$\mathbf{E}(|U_1|) = \sum_{v \in V} \mathbf{Pr}((v \notin \mathcal{T}) \wedge \mathcal{A}_{t_1}(v)) \geq -T + \sum_{v \in V} \mathbf{Pr}(\mathcal{A}_{t_1}(v)). \quad (30)$$

Here we subtract T to account for visits before the mixing time T .

Applying Lemma 4 we see that

$$\begin{aligned} \mathbf{E}(|U_1|) &\geq -T + \sum_{v \in V} \left(\frac{1 + \xi_{v,2}}{(1 + p_v)^{t_1}} + o(Te^{-\lambda t_1/2}) \right) \\ &= -T + (1 - o(1)) \sum_{v \in V} e^{-\pi_v(1-\epsilon_1)t^*} \\ &\geq -T + (1 - o(1)) e^{\epsilon_1 \theta t^*/n} \sum_{v \in V} e^{-\pi_v t^*} \\ &\geq -T + (1 - o(1)) n^{\epsilon_1 \theta} \\ &\approx n^{\epsilon_1 \theta} \rightarrow \infty, \end{aligned}$$

after using (22) to lower bound $e^{\epsilon_1 \theta t^*}$ and (19) to bound T , (here $\psi = o(\epsilon_1)$).

We summarise this as

$$\mathbf{E}(|U_1|) \approx \sum_{v \in V} \mathbf{Pr}(\mathcal{A}_{t_1}(v)) \approx n^{\epsilon_1 \theta}. \quad (31)$$

We now use the second moment method to show that $|U_1| > 0$ w.h.p. Fix two vertices v, w distinct from the start u of the walk. Let $\Gamma = \Gamma_{v,w}$ be obtained from G by identifying v, w as a single vertex $\gamma = \gamma_{v,w}$ and keeping the loop if $\{v, w\} \in E(G)$.

There is a natural measure preserving map from the set of walks in G which start at u and do not visit v or w , to the corresponding set of walks in Γ which do not visit γ . Thus the probability that $\mathcal{W}_u$ does not visit v or w in the first t steps is equal to the probability that a random walk $\widehat{\mathcal{W}}_u$ in Γ which also starts at u does not visit γ in the first t steps.

We first check that Lemma 4 can be applied to $\widehat{\mathcal{W}}_u$. We observe that it is valid to use T as a mixing time. This follows from Corollary 3.27 of [1] viz. that the relaxation time of a collapsed chain is bounded from above by that of the uncollapsed chain. Our estimate for R_γ should now be $1 + O(2T/(\theta n))$ (the 2 coming from vertices that are neighbors of v and w in G). Now

$$\frac{\pi_\gamma}{R_\gamma} = \frac{\pi_v + \pi_w}{1 + O(T/(\theta n))} = \left(1 + O\left(\frac{T}{\theta n}\right) \right) \left(\frac{\pi_v}{R_v} + \frac{\pi_w}{R_w} \right).$$

And so since $t_1 = \Theta(n \log n)$ we have

$$\begin{aligned} \Pr(v, w \in U_1) &= \Pr((v, w \notin \mathcal{T}) \wedge \mathcal{A}_{t_1}(v) \wedge \mathcal{A}_{t_1}(w)) = \Pr((\gamma \notin \mathcal{T}) \wedge \mathcal{A}_{t_1}(\gamma)) \leq \Pr(\mathcal{A}_{t_1}(\gamma)) \\ &= (1 + O(\omega^{-1})) \exp \left\{ -\frac{\pi_\gamma}{R_\gamma} t_1 \right\} = (1 + O(\omega^{-1})) \exp \left\{ -\left(\frac{\pi_v}{R_v} + \frac{p_w}{R_w} \right) t_1 \right\} \\ &= (1 + O(\omega^{-1})) \Pr(\mathcal{A}_{t_1}(v)) \Pr(\mathcal{A}_{t_1}(w)). \end{aligned}$$

It follows therefore that after using (31),

$$\mathbf{E}(|U_1|^2) \leq \mathbf{E}(|U_1|) + (1 + O(\omega^{-1})) \mathbf{E}(|U_1|)^2.$$

So, by the Chebyshev inequality,

$$\Pr(|U_1| = 0) \leq \frac{\mathbf{E}(|U_1|^2) - \mathbf{E}(|U_1|)^2}{\mathbf{E}(|U_1|)^2} \leq \frac{1}{\mathbf{E}(|U_1|)} + \frac{1}{\omega} \leq \frac{2}{n^{\epsilon_1 \theta}}. \quad (32)$$

This implies that $C_G \geq (1 - o(1))t_1$ and completes the proof of Theorem 1.

We will need the following lemma in Section 5.2. Let

$$\epsilon_2 = \frac{1}{\log^{1/4} n}.$$

Lemma 6.

$$\Pr(|T_{cov}(u) - t^*| \geq \epsilon_2 t^*) \leq 3\epsilon_2.$$

Proof. The probabilistic lower bound for $T_{cov}(u)$ follows from (32). For the upper bound, for a given $\alpha > 0$, we let $P_\alpha = \Pr(T_{cov}(u) \geq (1 + \alpha)t^*)$ and then we have for some large constant $K > 0$ that from (28),

$$\begin{aligned} \left(1 + \frac{2}{\theta \log n}\right) t^* &\geq \mathbf{E}(T_{cov}(u)) \geq \\ &\mathbf{E}(T_{cov}(u) \mid T_{cov}(u) \leq t_1) \Pr(T_{cov}(u) \leq t_1) + \\ &\mathbf{E}(T_{cov}(u) \mid t_1 < T_{cov}(u) \leq (1 + \alpha)t^*) \Pr(t_1 < T_{cov}(u) \leq (1 + \alpha)t^*) \\ &+ \mathbf{E}(T_{cov}(u) \mid (1 + \alpha)t^* < T_{cov}(u) \leq Kt^*) \Pr((1 + \alpha)t^* < T_{cov}(u) \leq Kt^*) \\ &\geq 0 + t_1 \left(1 - \frac{2}{n^{\epsilon_1 \theta}} - P_\alpha\right) + (1 + \alpha)t^* (P_\alpha - O(n^{-L})). \quad (33) \end{aligned}$$

Here K and L are related as in Remark 5. We obtain

$$\Pr(t_1 < T_{cov}(u) \leq (1 + \alpha)t^*) \geq \left(1 - \frac{2}{n^{\epsilon_1 \theta}} - P_\alpha\right)$$

from (32).

It follows from (33), after division by t^* , that

$$1 + \frac{2}{\theta \log n} \geq 0 + \left(1 - \frac{2}{n^{\epsilon_1 \theta}} - \frac{2}{\log^{1/2} n} - P_\alpha\right) + (1 + \alpha)P_\alpha.$$

We deduce from this that $\alpha P_\alpha \leq \frac{3}{\log^{1/2} n}$ and then that $P_\alpha \leq \frac{3}{\log^{1/4} n}$ for $\alpha = \frac{1}{\log^{1/4} n}$. $\square$

4 Partitioning the graph

Notation: For sets $S \subseteq X \subseteq V$, let $\deg_X(v)$ denote the number of neighbors of v in X , and let $\deg_X(S) = \sum_{v \in S} \deg_X(v)$. We will reserve the un-subscripted $\deg$ for $\deg_V$. For given $S \subseteq X \subseteq V$, we also use X as the subgraph $G[X]$ of G induced by X in the notation $\Phi_X, \Phi_X(S)$.

We assume that the minimum degree $\delta(G) \geq \theta n$ for some constant $\theta > 0$ and that $\psi = 1/\log^{2/3} n$ as in (4). Suppose that $\zeta = n^{-\theta\psi}$.

We partition V as follows: our initial partition Π_0 consists of V alone. Suppose that we have created a partition Π , and $X \in \Pi$. We can use the algorithm of Leighton and Rao, [19], to find a cut $(S : \bar{S})$ of X such that $\Phi_X \leq \Phi_X(S) \leq c_{LR} \Phi_X \log n$, where $c_{LR} > 0$ constant. If $\Phi_X(S) \geq \zeta$, we do not partition X any further. Otherwise, if $\Phi_X(S) < \zeta$, we refine Π by splitting X into $X_1 = S$ and $X_2 = X \setminus S$. For $\ell = 1, 2$ let

$$Y_\ell = Y(X_\ell) = \left\{v \in X_\ell : \deg_{X_\ell}(v) \leq \deg_{X_{3-\ell}}(v)\right\}. \quad (34)$$

We replace X in the partition Π by the pair

$$Z_\ell = (X_\ell \cup Y_{3-\ell}) \setminus Y_\ell \text{ for } \ell = 1, 2. \quad (35)$$

Suppose that, for all $v \in X$, $\deg_X(v) \geq \beta n$, where (see (37)) $\beta = \min_d \beta_d$ satisfies $\beta > 4\zeta^{1/2}$, for $d = O(1)$. If $v \in Y_\ell$, $\deg_{X \setminus X_\ell}(v) \geq \beta n/2$, and thus

$$|Y_\ell| \leq \frac{e_X(X_\ell, X \setminus X_\ell)}{\beta n/2} \leq \frac{2\zeta n}{\beta} \leq \frac{\zeta^{1/2} n}{2}. \quad (36)$$

For the second inequality we used the crude bound, $e_X(X_\ell, X \setminus X_\ell) \leq \zeta n^2$, which follows from $\Phi_X(S) < \zeta$ and (1).

Continue in this way until the output of the algorithm of [19] returns a cut $(S, V_i \setminus S)$ such that $\Phi_{V_i}(S) \geq \zeta$ for all sets of the partition $\Pi = (V_1, V_2, \dots, V_s)$. The depth $d_\Pi(V_i)$ of V_i in Π is defined as follows: $d_{\Pi_0}(V) = 0$ and if $X \in \Pi$ has depth d , then its *descendants* Z_1, Z_2 will both have depth $d + 1$. Suppose that V_i has depth d . We claim that $d = O(1)$ and that

$$\min \{\deg_{V_i}(v) : v \in V_i\} \geq \beta_d n, \quad \text{where } \beta_d = \frac{\theta}{3^d}. \quad (37)$$

If so it follows from (37), and $d = O(1)$, that at depth d , $|V_i| \geq \beta_d n = \Omega(n)$.

We prove (37) by induction. It is true for $d = 0$ and $\beta_0 = \theta$. If V_i has depth $d + 1$ and arises from splitting X at depth d then for $v \in V_i$, provided $d = O(1)$, then from (36)

$$\deg_{V_i}(v) \geq \frac{\beta_d n}{2} - |Y(X_i)| \geq \left(\frac{\beta_d}{2} - \frac{2\zeta}{\beta_d} \right) n \geq \beta_{d+1} n. \quad (38)$$

We also have

$$|\{v \in V_i : \deg_{V_i}(v) \leq \deg(v) - d\zeta^{1/2}n\}| \leq 3d\zeta^{1/2}n. \quad (39)$$

This follows by induction. It is true for $d = 0$. If V_i has depth $d + 1$ and arises from splitting X at depth d then

$$\begin{aligned} & |\{v \in V_i : \deg_{V_i}(v) \leq \deg(v) - (d+1)\zeta^{1/2}n\}| \leq \\ & |\{x \in X : \deg(x) \leq \deg(v) - d\zeta^{1/2}n\}| + 2\zeta^{1/2}n + \zeta^{1/2}n. \end{aligned}$$

The first term on the RHS is the number of vertices which have low degree at the previous level. The next term counts the at most $2\zeta^{1/2}n$ vertices which lose at least $\zeta^{1/2}n/2$ edges, as the cut (X_1, X_2) of X which gave rise to (V_1, V_2) has at most ζn^2 edges. The last term comes from (36) and compensates for the neighbours of $Y(X_\ell)$ with at most $\zeta^{1/2}n/2$ edges in the cut, who lost degree (at most) $|Y(X_\ell)| \leq \zeta^{1/2}n/2$, when $Y(X_\ell)$ was moved out of X_ℓ to obtain V_i .

It follows from (37) and (39) that if $d = O(1)$ and $X = V_1 \cup V_2$ then,

$$|V_1|, |V_2| \geq \theta n/2. \quad (40)$$

To see this, suppose that X was initially partitioned into $X_1 = S, X_2 = X \setminus S$ where $\Phi_X(S) < \zeta$ and that $V_i = (X_\ell \cup Y_{3-\ell}) \setminus Y_\ell$, as in (34), (35) above, with V_i replacing $Z_i, i = 1, 2$. As $Y_1 = X_1 \setminus V_1$ and $Y_2 = V_1 \setminus X_1$, (36) implies that $|(V_1 \setminus X_1) \cup (X_1 \setminus V_1)| = o(n)$. This and (37) implies that $|X_1| = \Omega(n)$. Suppose that $|V_1| < \theta n/2$. Then (39) implies that there are $|X_1|(1 - o(1))$ vertices of X_1 of degree at least $\theta n(1 - o(1))$ in X . So there are at least $\theta |X_1| n(1 - o(1))/2$ edges in the cut $X_1 : X_2$. And hence,

$$\zeta > \Phi_X(X_1) \geq \frac{(\theta |X_1| n(1 - o(1))/2) \times d(V)}{|X_1| n \times d(V)} \geq \frac{\theta}{3},$$

which is a contradiction.

By (40) we have that sets at depth d have size at least $\theta n/2$. On the other hand, at least $\theta n/2$ vertices are moved at each partition step, and so sets at depth d have size at most $n - d\theta n/2$. This means that $n - d\theta n/2 \geq \theta n/2$, and partitioning must stop when $d < 2/\theta$.

5 Computing the cover time

Let $V_1, V_2, \dots, V_s$ be as in Section 4. For each i we add weighted edges to create a multi-graph H_i such that a random walk on H_i corresponds to the visits to V_i of a random walk on G . Thus, for each i , we define H_i by adding extra edges to $E(V_i)$. If $v, w \in V_i$ then we add an *oriented* edge (v, w) and give it a weight $\rho_{v,w}$. Here $\rho_{v,w}$ is the probability that a walk started at v leaves V_i immediately and returns to V_i at w and we have $\sum_{w \in V_i} \rho_{v,w} = \frac{\deg(v) - \deg_{V_i}(v)}{\deg(v)}$. The (unoriented) edges of G contained in V_i will be given weight one. We will use $w()$ to denote weight in H_i .

Remark 6. If we take the random walk $\mathcal{W}_u = (u = X(0), X(1), \dots, X(t), \dots)$ and delete the entries $X(t)$ that are not in V_i then the remaining sequence is a random walk $\mathcal{Z}_i$ on H_i .

A random walk $\mathcal{Z}_i$ on H_i will have steady state $\pi_{v,i} = \deg(v)/\deg(V_i)$, $v \in V_i$ and will satisfy the conditions of Theorem 1. Indeed, the walk is reversible. Checking detailed balance, we have

$$\begin{aligned} \pi_{v,i} P_i(v, w) &= \frac{\deg(v)}{\deg(V_i)} \left(\frac{\deg_{V_i}(v)}{\deg(v)} \cdot \frac{1_{\{v,w\} \in E(H)}}{\deg_{V_i}(v)} + \rho_{v,w} \right) = \\ &= \frac{1_{\{v,w\} \in E(H)} + \deg(v) \rho_{v,w}}{\deg(V_i)} = \pi_{w,i} P_i(w, v), \end{aligned}$$

since necessarily, $\deg(v) \rho_{v,w} = \deg(w) \rho_{w,v}$. This follows from the fact that for any individual walk $W = (x_0 = v, x_1, x_2, \dots, x_k = w)$ from v to w on G and its reversal $\bar{W} = (x_k = w, x_{k-1}, \dots, x_1, x_0 = v)$ we have

$$\pi_{v,i} \mathbf{Pr}(W) = \frac{\deg(v)}{\deg(V_i)} \prod_{j=0}^{k-1} \frac{1}{\deg(x_j)} = \frac{\deg(w)}{\deg(V_i)} \prod_{j=1}^k \frac{1}{\deg(x_j)} = \pi_{w,i} \mathbf{Pr}(\bar{W}).$$

To obtain $\deg(v) \rho_{v,w} = \deg(w) \rho_{w,v}$ we sum over all walks from v to w with interior vertices not in V_i .

Now consider the conductance of H_i . In what follows we use the fact that the weight of edges incident with a vertex v in H_i is equal to the degree of v in G . Suppose that $S \subseteq V_i$. According to the definition of $\Phi_{H_i}(S)$ (or rather its extension to graphs with weighted edges),

$$\Phi_{H_i}(S) = \frac{w(S, \bar{S}) w(V_i)}{\deg(S) \deg(\bar{S})} \geq \frac{e(S, \bar{S}) \deg(V_i)}{|S| |\bar{S}| n^2} = \Phi_G(S) \frac{\deg(V_i)}{n^2} \geq \Phi_G(S) \frac{\theta^2}{3},$$

assuming $|V_i|$ has at least $\theta n/2$ vertices of degree $\theta n(1 - o(1))$. Thus

$$\Phi(H_i) \geq \Phi(G) \frac{\theta^2}{3} \geq \frac{\Phi_G(S)}{c_{LR} \log n} \frac{\theta^2}{3} \geq \frac{\zeta \theta^2}{3 c_{LR} \log n},$$

where the last step comes from combining Theorem 2 of Leighton and Rao [19], with equation (3) of that paper, that gives a deterministic polynomial algorithm to find a cut $(S : \bar{S})$ such that $\Phi(G) \leq \Phi_G(S) \leq c_{LR} \Phi(G) \log n$.

5.1 When $T = o(C)$: Proof of Theorem 2

We see from Theorem 1 that $\mathcal{Z}_i$ will have to make a number of steps in some explicit range $(1 \pm \epsilon)C_i$ in order to cover V_i .

Remark 7. *The estimate C_i does not depend significantly on the values $\rho_{v,w}$. We see from Theorem 1 that up to a factor $(1 + o(1))$, the C_i depend only on the degrees of H_i . But we can compute close approximations to the $\rho_{v,w}$. For this we need to compute the values*

$$\sigma_{x,y,i,t} = \Pr(\mathcal{W}_x(t) = y, \mathcal{W}_x(\tau) \notin V_i, 1 \leq \tau \leq t). \quad (41)$$

Given these values, we have

$$\rho_{v,w} = \sum_{x,y \notin V_i} P(v,x)P(y,w) \sum_{t \geq 1} \sigma_{x,y,i,t}.$$

Finally, to compute the values in (41) we simply look at powers of the matrix Q_i that is obtained from P by replacing entries in columns associated with V_i by zeroes.

Consider a walk $\mathcal{W}$ starting in the steady state that walks for t steps. The expected number of visits to V_i is $t\pi(V_i)$ and it will be concentrated around this, if the mixing time of $\mathcal{W}$ is small. For example, Corollary 2.1 of Paulin [22] shows that if $Z_{i,t}$ is the number of visits to V_i then

$$\Pr(|Z_{i,t} - t\pi(V_i)| \geq u) \leq \exp \left\{ -\frac{2u^2}{tT} \right\}. \quad (42)$$

Next let

$$C = \max \left\{ \frac{C_i}{\pi(V_i)} : i \in [s] \right\} = \Omega(n \log n). \quad (43)$$

Then $C_G = (1 \pm \epsilon)C$ if $T = o(C)$. Indeed, putting $t = C, u = \epsilon C$ in (42) we see immediately that w.h.p. $T_{cov}(u)$ is within a factor $1 + o(1)$ of C . This immediately gives us a lower bound of $(1 - o(1))C$ for the expectation. For the upper bound we use Remark 5 in the following way: we know that $\Pr(T_{cov}(u) \in [(1 + o(1))C, KC]) = o(1)$ and so this range adds $o(C)$ to the expectation. After this, $[KC, \infty]$ adds a negligible amount for large K .

This completes our proof of Theorem 2.

5.2 When T is large: Proof of Theorem 3

While a nice formula for the cover time is not necessarily attainable, we claim that we can deterministically compute quantities that give us a factor $2 + o(1)$ estimate for the cover time, in a time polynomial in n .

Consider the $n \times n$ matrix Q where $Q(u, v) = P(u, v)\xi_{\phi(u)}$ where the $\xi_i, i \in [s]$ are indeterminate and ϕ is defined by $u \in V_{\phi(u)}$, for $V_i \in \Pi$. Now consider the t -th power of Q . Then

$$Q^t(u, v) = \sum_{\tau_1 + \dots + \tau_s = t} \Pr(\mathcal{W} \text{ goes from } u \text{ to } v \text{ in } t \text{ steps and makes } \tau_i \text{ } V_i\text{-moves, for } i \in [s]) \prod_{i=1}^s \xi_i^{\tau_i}.$$

Here a V_i -move is from a vertex in V_i to any vertex $v \in V$. Note that the number of V_i moves is equal to the number of moves by $\mathcal{Z}_i$. The cover time of any connected n -vertex graph of minimum degree δ is $O(n|E|/\delta)$, [17]. When $\delta = \theta n$, $C_G = O(n^2)$. Thus we compute Q^t for $1 \leq t \leq n^4$, and observe that this computation can be done in $O(n^7)$ time. Let

$\kappa(u, \tau, i)$ denote the number of steps in $\mathcal{W}_u$ needed for τ V_i -moves.

Next let $C_i^\pm = (1 \pm \epsilon_2) C_i$ be such that the cover time of $\mathcal{Z}_i$ is in $[C_i^-, C_i^+]$ w.h.p., see Lemma 6.

Note that the C_i are given by (5) of Theorem 1, which can be computed in deterministic polynomial time.

Let $U_{i,t}$ denote the set of unvisited vertices of V_i at time t . We know from the proof of Theorem 1 that w.h.p. if $t \leq \kappa(u, C_i^-, i)$ then $U_{i,t} \neq \emptyset$. This implies that

$$C_G \geq \max_{u \in V} \mathbf{E}(\max_{i \in [s]} \kappa(u, C_i^-, i)). \quad (44)$$

For the RHS of (9), we note that at time $\max_{i \in [s]} \kappa(u, C_i^-, i)$ the walk $\mathcal{W}_u$ will be at some vertex v and then after a further $\max_{i \in [s]} \kappa(v, C_i^-, i)$ steps¹ the walk $\mathcal{W}_u$ will w.h.p. have spent at least time $2C_i^-$ in V_i for every $i \in [s]$.

Because $2C_i^- > C_i^+$, the walk $\mathcal{W}_u$ will w.h.p. have covered V . Thus

$$C_G \leq (2 + o(1)) \max_{u \in V} \mathbf{E}(\max_{i \in [s]} \kappa(u, C_i^-, i)). \quad (45)$$

and this completes the proof of Theorem 3.

¹We could write $\max_{i \in [s]} \kappa(v, C_i^+ - C_i^-, i)$ here, but we cannot prove that this is significantly smaller than what we have written.

References

- [1] D. Aldous and J. Fill. Reversible Markov Chains and Random Walks on Graphs, <http://stat-www.berkeley.edu/pub/users/aldous/RWG/book.html>.
- [2] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász and C. Rackoff, Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems. *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science* (1979) 218-223.
- [3] J. Brown and R. Churchill, *Complex Variables and Applications*, (Sixth Edition) McGraw-Hill (1996).
- [4] A.K. Chandra, P. Raghavan, W.L. Ruzzo and R. Smolensky, The electrical resistance of a graph captures its commute and cover times, *Proceedings of 21st. ACM Symposium on Theory of Computing*, (1989) 574-586.
- [5] C. Cooper and A. M. Frieze, The cover time of random regular graphs, *SIAM Journal on Discrete Mathematics*, 18 (2005) 728-740.
- [6] C. Cooper and A. M. Frieze, The cover time of the preferential attachment graph. *Journal of Combinatorial Theory Series B*, Series B, 97(2) (2007) 269-290.
- [7] C. Cooper and A. M. Frieze, The cover time of sparse random graphs, *Random Structures and Algorithms* 30 (2007) 1-16.
- [8] C. Cooper and A. M. Frieze, The cover time of the giant component of a random graph, *Random Structures and Algorithms*, 32, 401-439 (2008).
- [9] C. Cooper and A. M. Frieze, The cover time of random digraphs, Stationary distribution and cover time of random walks on random digraphs, *Journal of Combinatorial Theory B* 102 (2012) 329-362.
- [10] C. Cooper and A. M. Frieze, The cover time of random geometric graphs, *Random Structures and Algorithms* 38 (2011) 324-349.
- [11] C. Cooper, A.M. Frieze and T. Radzik, The cover time of random walks on random uniform hypergraphs, *Proceedings of SIROCCO 2011*, 210-221.
- [12] J. Ding, Asymptotic of cover times via Gaussian free fields: Bounded degree graphs and general trees, *The Annals of Probability* 42 (2014) 464-496.
- [13] J. Ding, J.R. Lee and Y. Peres, Cover times, blanket times, and majorizing measures, *Annals of Mathematics* 175 (2012) 1409-1471.
- [14] U. Feige, A tight upper bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 51-54.

- [15] U. Feige, A tight lower bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 433-438.
- [16] J. Kahn, J.H. Kim, L. Lovász and V. Vu, The cover time, the blanket time, and the Matthews bound, in *41st Annual IEEE Symposium on Foundations of Computing* (Redondo Beach, CA) (2000).
- [17] J. Kahn, N. Linial, N. Nisan, M. Saks, On the Cover Time of Random Walks on Graphs, *Journal of Theoretical Probability*, Vol. 2, No. 1, (1989).
- [18] D. Levin, Y. Peres and E. Wilmer, *Markov Chains and Mixing Times: Second Edition*, American Mathematical Society, (2017).
- [19] T. Leighton and S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, *Journal of the ACM* 46 (1999) 787-832
- [20] P. Matthews, Covering problems for Brownian motion on spheres, *Annals of Probability* 16 (1988) 189-199.
- [21] R. Meka, A polynomial time approximation scheme for computing the supremum of Gaussian processes, *Annals of Applied Probability* 25 (2015) 465-476.
- [22] D. Paulin, Concentration inequalities for Markov chains by Marton couplings and spectral methods, *Electronic Journal of Probability*, (2015).
- [23] H. Wilf, *Generatingfunctionology*, Academic Press (1990).
- [24] A. Zhai, Exponential concentration of cover times. *Electron. J. Probab.*, Volume 23 (2018), paper no. 32, 22 pp.

A Proof of Lemma 4

Write

$$R(z) = R_T(z) + \widehat{R}_T(z) + \frac{\pi_v z^T}{1-z}, \quad (46)$$

where $R_T(z)$ is given by (13) and

$$\widehat{R}_T(z) = \sum_{t \geq T} (r_t - \pi_v) z^t$$

generates the error in using the stationary distribution π_v for r_t when $t \geq T$. Similarly,

$$H(z) = \widehat{H}_T(z) + \frac{\pi_v z^T}{1-z}. \quad (47)$$

Equation (10) implies that the radii of convergence of both $\widehat{R}_T$ and $\widehat{H}_T$ exceed $1+2\lambda$. Moreover, for $Z = H, R$ and $|z| \leq 1 + \lambda$, we see from (10) that

$$|\widehat{Z}_T(z)| \leq \frac{\pi_v}{2} \sum_{t \geq T} \left(\frac{2(1+\lambda)}{\omega} \right)^{\lfloor t/T \rfloor} \leq \frac{2(1+\lambda)T\pi_v}{\omega} = O(\omega^{-2}). \quad (48)$$

Using (46), (47) we rewrite $F(z) = H(z)/R(z)$ from (12) as $F(z) = B(z)/A(z)$ where

$$A(z) = \pi_v z^T + (1-z)(R_T(z) + \widehat{R}_T(z)), \quad (49)$$

$$B(z) = \pi_v z^T + (1-z)\widehat{H}_T(z). \quad (50)$$

For real $z \geq 1$ and $Z = H, R$, we have

$$Z_T(1) \leq Z_T(z) \leq Z_T(1)z^T.$$

Let $z = 1 + \beta\pi_v$, where $0 \leq \beta \leq 1$. Since $T\pi_v \leq \omega^{-1}$ we have

$$Z_T(z) = Z_T(1)(1 + \xi_1) \text{ where } |\xi_1| \leq (1 + \beta\pi_v)^T - 1 \leq \frac{2\beta}{\omega}.$$

$T\pi_v \leq \omega^{-1}$ and $R_v \geq 1$ implies that

$$A(z) = \pi_v(1 - \beta R_v(1 + \xi_1)) \text{ where } |\xi_1| = O(\omega^{-1}).$$

It follows that $A(z)$ has a real zero at z_0 , where

$$z_0 = 1 + \frac{\pi_v}{R_v(1 + \xi_1)} = 1 + p_v. \quad (51)$$

We also see that since $|z_0^T| \leq 1 + 2\omega^{-1}$,

$$\begin{aligned} A'(z_0) &= T\pi_v z_0^{T-1} - (R_T(z_0) + \widehat{R}_T(z_0)) - p_v(R'_T(z_0) + \widehat{R}'_T(z_0)) \\ &= O(\omega^{-1}) - (R_v + O(\omega^{-1}) + o(\omega^{-1})) - o(\pi_v) \\ &= -R_v + O(\omega^{-1}) \\ &\neq 0. \end{aligned}$$

and thus z_0 is a simple zero (see e.g. [3] p193). The value of $B(z)$ at z_0 is

$$B(z_0) = \pi_v (1 + O(\omega^{-1}) + o(\omega^{-1})) = \pi_v (1 + O(\omega^{-1})) \neq 0. \quad (52)$$

Thus,

$$\frac{B(z_0)}{A'(z_0)} = -(1 + \xi_2) p_v \text{ where } |\xi_2| = O(\omega^{-1}). \quad (53)$$

Thus (see e.g. [3] p195) the principal part of the Laurent expansion of $F(z)$ at z_0 is

$$f(z) = \frac{B(z_0)/A'(z_0)}{z - z_0}. \quad (54)$$

To approximate the coefficients of the generating function $F(z)$, we now use a standard technique for the asymptotic expansion of power series (see e.g.[23] Theorem 5.2.1).

We prove below that $F(z) = f(z) + g(z)$, where $g(z)$ is analytic in $C_\lambda = \{|z| \leq 1 + \lambda\}$ and that

$$M = \max_{z \in C_\lambda} |g(z)| = O(\omega^{-1}).$$

Let $a_t = [z^t]g(z)$, then (see e.g.[3] p143), $a_t = g^{(t)}(0)/t!$. By the Cauchy Inequality (see e.g. [3] p130) we see that $|g^{(t)}(0)| \leq Mt!/(1 + \lambda)^t$ and thus

$$|a_t| \leq \frac{M}{(1 + \lambda)^t} \leq Me^{-t\lambda/2}.$$

As $[z^t]F(z) = [z^t]f(z) + [z^t]g(z)$ and $[z^t]1/(z - z_0) = -1/z_0^{t+1}$ we have

$$[z^t]F(z) = \frac{-B(z_0)/A'(z_0)}{z_0^{t+1}} + \eta_1(t) \text{ where } |\eta_1(t)| \leq Me^{-t\lambda/2}. \quad (55)$$

Thus, we obtain

$$[z^t]F(z) = \frac{(1 + \xi_2)p_v}{(1 + p_v)^{t+1}} + \eta_1(t).$$

Now

$$\Pr(\mathbf{A}_t(v)) = \sum_{\tau > t} f_\tau(u \rightarrow v) = \sum_{\tau > t} \left(\frac{(1 + \xi_2)p_v}{(1 + p_v)^{\tau+1}} + \eta_1(\tau) \right) = \frac{1 + \xi_2}{(1 + p_v)^{t+1}} + \eta_2(t),$$

where

$$\eta_2(t) = \sum_{\tau > t} \eta_1(\tau) \leq \frac{Me^{-\lambda t/2}}{1 - e^{-\lambda/2}} = o(Te^{-\lambda/2}).$$

This completes the proof of (17).

Now $M = \max_{z \in C_\lambda} |g(z)| \leq \max |f(z)| + \max |F(z)| = O(T\pi_v) + \max |F(z)| = O(\omega^{-1}) + \max |F(z)|$, where $F(z) = B(z)/A(z)$. On C_λ we have, using (48)-(50),

$$|F(z)| \leq \frac{\pi_v z^T + o(\pi_v)}{\pi_v z^T + \lambda(|R_T(z)| - O(\omega^{-2}))} = O\left(\frac{\pi_v z^T}{T^{-1}R_v}\right) = O(\omega^{-1}).$$

We now prove that z_0 is the only zero of $A(z)$ inside the circle C_λ and this implies that $F(z) - f(z)$ is analytic inside C_λ . We use Rouché's Theorem (see e.g. [3]), the statement of which is as follows: *Let two functions $\phi(z)$ and $\gamma(z)$ be analytic inside and on a simple closed contour C . Suppose that $|\phi(z)| > |\gamma(z)|$ at each point of C , then $\phi(z)$ and $\phi(z) + \gamma(z)$ have the same number of zeroes, counting multiplicities, inside C .*

Let the functions $\phi(z), \gamma(z)$ be given by $\phi(z) = (1 - z)R_T(z)$ and $\gamma(z) = \pi_v z^T + (1 - z)\widehat{R}_T(z)$.

$$\frac{|\gamma(z)|}{|\phi(z)|} \leq \frac{\pi_v(1 + \lambda)^T}{\lambda\theta} + \frac{|\widehat{R}_T(z)|}{\theta} = o(1).$$

As $\phi(z) + \gamma(z) = A(z)$ we conclude that $A(z)$ has only one zero inside the circle C_λ . This is the simple zero at z_0 .