Optimal Construction of Edge-Disjoint Paths in Random Regular Graphs

Alan M. Frieze[*]   Lei Zhao[†]

Abstract

Given a graph $G = (V, E)$ and a set of κ pairs of vertices in V, we are interested in finding for each pair (a_i, b_i), a path connecting a_i to b_i, such that the set of κ paths so found is edge-disjoint. (For arbitrary graphs the problem is $\mathcal{NP}$-complete, although it is in $\mathcal{P}$ if κ is fixed.)

We present a polynomial time randomized algorithm for finding the optimal number of edge disjoint paths (up to constant factors) in the random regular graph $G_{n,r}$, for r sufficiently large. (The graph is chosen first, then an adversary chooses the pairs of endpoints.)

1 Introduction

Given a graph $G = (V, E)$ with n vertices, and a set of κ pairs of vertices in V, we are interested in finding for each pair (a_i, b_i), a path connecting a_i to b_i, such that the set of κ paths so found is edge-disjoint.

For arbitrary graphs the related decision problem is $\mathcal{NP}$-complete, although it is in $\mathcal{P}$ if κ is fixed – Robertson and Seymour [23]. Nevertheless, this negative result can be circumvented for certain classes of graphs. Peleg and Upfal [22] presented a polynomial time algorithm for the case where G is a (sufficiently strong) bounded degree expander graph, and $\kappa \leq n^\epsilon$ for a small constant ϵ that depends on the expansion property of the graph. (A precise upper bound for ϵ was not computed, but it is clearly less that $1/3$). This result has been improved and extended by Broder, Frieze, and Upfal [10, 11], Frieze [13], Leighton and Rao [18] and Leighton, Rao and Srinivasan [19, 20]: G still has to be a (sufficiently strong) bounded degree expander but κ can now grow as fast as $n/(\ln n)^\theta$, where θ depends only on the expansion properties of the input graph, but is at least 2.

In random graphs Shamir and Upfal have shown in [24] that any set of up to $O(\sqrt{n})$ pairs can be connected via vertex-disjoint paths; similar results using efficient flow techniques were also obtained by Hochbaum [15]. These results were proved for graphs with $m \geq Kn \log n$ random edges, where K is a sufficiently large constant. In these two papers the pairs are chosen *first* and then the graph is randomly generated. Thus they do not deal with the problem of satisfying *all* sets of κ pairs.

Let D be the median distance between pairs of vertices in G. Clearly it is not possible to connect more than $O(m/D)$ pairs of vertices by edge-disjoint paths, for all choices of pairs,

since some choice would require more edges than all the edges available. In the case of bounded degree expanders, this absolute upper bound on κ is $O(n/\log n)$. The results mentioned so far use only a vanishing fraction of the set of edges of the graph, thus are far from reaching this upper bound. Broder, Frieze, Suen and Upfal [8] have proved that for the model of random graphs, $G_{n,m}$, the absolute upper bound is achievable to within a constant factor, but only if the average degree is at least $\ln n$. In this work, we show that this result holds when the minimum degree is a large enough constant. Without significant loss of generality, we consider random regular graphs and present an algorithm that constructs the required paths in polynomial time.

As usual, let $G_{n,p}$ denote a random graph with vertex set $\{1, 2, \ldots, n\} = [n]$ in which each possible edge is included independently with probability p, and let $G_{n,m}$ denote a random graph also with vertex set $[n]$ and exactly m edges, all sets of m edges having equal probability. Let $G_{n,r}$ denote a random graph again with vertex set $[n]$ which is chosen uniformly at random from all possible r-regular graphs.

Our main result is formulated in the following theorem.

Theorem 1 *Let r be a sufficiently large constant. Then, as $n \longrightarrow \infty$, the graph $G_{n,r}$ has the following property* **whp**: *there exist positive absolute constants α, β such that for all sets of pairs of vertices $\{(a_i, b_i) \mid i = 1, ..., \kappa\}$ satisfying:*

(i) $\kappa = \lceil \alpha r n / \log_r n \rceil$,
(ii) for each vertex v, $|\{i : a_i = v\}| + |\{i : b_i = v\}| \leq \beta r$,

there exist edge-disjoint paths in G, joining a_i to b_i, for each $i = 1, 2, \ldots, \kappa$. Furthermore, there is an $O(n^3)$ time randomized algorithm for constructing these paths.

This result is the best possible up to constant factors. The median distance between pairs of vertices in G is $\Omega(\log_r n)$. The need for $\beta < 1$ is discussed in [8].

The analogous problem of finding vertex disjoint paths in random graphs is dealt with in [9].

2 Preliminaries

The paper contains a few unspecified absolute constants of which α above is the first. Exact values could be given but it is easier for us *and the reader* if we simply give the relations between them. New constants will be introduced as $C_0, \ldots$ without further comment. Furthermore, specific constants have been chosen for convenience. We made no attempt to optimize them, and, in general, we only claim that inequalities dependent on n or r hold for n or r sufficiently large.

For a graph $G = (V, E)$ and $v \in V$ we let $d_G(v)$ denote the degree of v in G. We use $\delta(G)$ and $\Delta(G)$ to denote the smallest and largest degrees respectively. For a set $S \subseteq V$ we let $\bar{S} = V \setminus S$ and define its neighbor set, $N_G(S)$, as

$$N_G(S) = \{v \in \bar{S} : \exists w \in S \text{ such that } \{v, w\} \in E\}.$$

For $S, T \subseteq V$ we let $e_G(S, T)$ denote the number of edges with one end in S and one end in T. For $S \subseteq V$, we use $G[S]$ to denote the subgraph of G induced by S.

The *Chernoff bounds* on the tails of the Binomial $\text{Bin}(n, \theta)$ that we use are

$$\mathbf{Pr}(\text{Bin}(n, \theta) \leq (1 - \epsilon)n\theta) \leq e^{-\epsilon^2 n\theta/2}, \tag{1}$$

$$\mathbf{Pr}(\text{Bin}(n, \theta) \geq (1 + \epsilon)n\theta) \leq e^{-\epsilon^2 n\theta/3}, \tag{2}$$

valid for $0 \leq \epsilon \leq 1$.

3 Overview of the algorithm

Our algorithm divides naturally into the four phases sketched below.

Phase 0: Partition G into six edge-disjoint graphs $G_i = (V_i, E_i)$, $1 \leq i \leq 6$. Phase 1 will use only the graph G_1; Phase 2 will use only the graphs G_2, G_3, G_4 and G_5; and Phase 3 will use only the graph G_6. The partition is such that $V_1 = V$ but $V_2 = \cdots = V_6 \subseteq V$ with $|V_2| \geq n - \epsilon n$, where $\epsilon = \epsilon(r)$ is a small constant.

Phase 1: Choose a random set $Z = \{z_1, \ldots, z_{2\kappa}\}$ of 2κ points in V_2. Connect the endpoints $\{(a_i, b_i) \mid i = 1, \ldots, \kappa\}$ to the newly chosen points in an arbitrary manner via edge-disjoint paths in G_1 using a flow algorithm. Let $\tilde{a}_i$ (resp. $\tilde{b}_i$) be the vertex connected to a_i (resp. b_i). The original problem is now reduced to finding edge-disjoint paths from $\tilde{a}_i$ to $\tilde{b}_i$ for each i. (This randomization was used in [10] and has its roots in Valiant's routing algorithm [26].)

Phase 2: We split this into parts (a),(b).
(a) For each z_i, $1 \leq i \leq 2\kappa$, we do a random walk of length

$$\tau_0 = \lceil 10 \log_r n \rceil. \tag{3}$$

starting at z_i in G_2. The terminating endpoint of this walk will be denoted by $\hat{a}_j$ if $z_i = \tilde{a}_j$ and by $\hat{b}_j$ if $z_i = \tilde{b}_j$. Ater we complete a walk we remove the edges from the corresponding graph.

(b) For each i in turn, we first do a random walk starting at $\hat{b}_i$ and terminating at b_i^* in G_4. Then we repeatedly do a certain type of random walk in G_3 starting from $\hat{a}_i$ until one of these walks ends at b_i^*. We keep the last walk as our path from $\hat{a}_i$ to b_i^* and remove all edges seen in these walks from G_3 and G_4 respectively.

Most pairs $(\tilde{a}_i, \tilde{b}_i)$ will be successfully connected in this phase. For such a pair, the final path from a_i to b_i is the concatenation of the paths (after removing cycles if necessary.) from a_i to $\tilde{a}_i$ and from b_i to $\tilde{b}_i$ found in Phase 1, the paths from $\tilde{a}_i$ to $\hat{a}_i$, from $\hat{a}_i$ to b_i^*, from b_i^* to $\hat{b}_i$ and from $\tilde{b}_i$ to $\hat{b}_i$ which are found in this stage. It is important in our analysis to ensure that random walks are done on graphs of high minimum degree. We use G_5 as a *backup* for ensuring that this is done.

Phase 3: At the end of Phase 2, **whp**, there will be at most $n/(\ln n)^{C_0}$ pairs $(\tilde{a}_i, \tilde{b}_i)$ which have not been joined by paths. We use the algorithm of [10] to join them by edge disjoint paths, using only the edges of G_6, and then construct the final paths from a_i to b_i as above.

To prove Theorem 1 it suffices to show that for almost every $G_{n,r}$:

- Phases 0 and 1 will succeed for *all* choices of $a_1, \ldots, b_\kappa$ and *almost every* choice of $z_1, \ldots, z_{2\kappa}$.
- Phases 2 and 3 are successful for *almost every* choice of $z_1, \ldots, z_{2\kappa}$ and *any* bijection $\{\tilde{a}_1, \ldots, \tilde{a}_\kappa, \tilde{b}_1, \ldots, \tilde{b}_k\} \longrightarrow \{z_1, \ldots, z_{2\kappa}\}$

Note that to prove these facts we have to consider only one experiment, namely choose $G_{n,r}$ at random, choose $a_1, \ldots, b_\kappa$ arbitrarily and then $z_1, \ldots, z_{2\kappa}$ at random. From this we can deduce that almost every $G_{n,r}$ is such that for *all* choices of $a_1, \ldots, b_\kappa$ and almost every choice of $z_1, \ldots, z_{2\kappa}$, we can find edge-disjoint paths $a_i - \tilde{a}_i - \hat{a}_i - b_i^* - \hat{b}_i - \tilde{b}_i - b_i$ for $1 \leq i \leq \kappa$.

4 Detailed description of the algorithm

The input to our algorithm is a random graph $G_{n,r}$ and a set of pairs of vertices $\{(a_i, b_i) \mid i = 1, \ldots, \kappa\}$ satisfying the premises of Theorem 1. The output is a set of κ edge-disjoint paths, $P_1, \ldots, P_\kappa$ such that P_i connects a_i to b_i.

4.1 Phase 0.

We start by partitioning G into six edge-disjoint graphs $G_i = (V_i, E_i)$, for $1 \leq i \leq 6$. Phase 1 will use only G_1; Phase 2 will use only G_2, G_3, G_4 and G_5; Phase 3 will use only G_6. The partition is such that $V_1 = V$ but $V_2 = \cdots = V_6 \subseteq V$ with $|V_2| \geq n - \epsilon n$, where $\epsilon = \epsilon(r)$ is a small constant.

In this construction, we use the notion of a k-core. The k-core of a collection of graphs $H_1, H_2, \ldots, H_l$ on the same vertex set V is the largest subset of V which induces a subgraph of minimum degree at least k in each of $H_1, H_2, \ldots, H_l$. It is unique and can be found by repeatedly removing vertices which have degree less than k in one of the l induced graphs. This continues until there are no vertices left or minimum degree k is achieved.

The algorithm below starts by constructing preliminary versions of these graphs, denoted G_i' for $1 \leq i \leq 6$. Then edges and vertices are deleted from $G_2', \ldots, G_6'$ in order to achieve certain minimum degree properties.

1.   **Algorithm** SPLIT
2.   **begin**
3.   Divide E into E_i', $1 \leq i \leq 6$ by placing each edge of E independently
            with probability $1/6$ into each of E_i' for $1 \leq i \leq 6$.
4.   **for** $1 \leq i \leq 6$ set $G_i' \leftarrow (V, E_i')$
5. $K \leftarrow \lceil r/7 \rceil - $core of G_1'
6.   **for** $2 \leq i \leq 6$ set $G_i \leftarrow G_i'[K]$
7. $K \leftarrow \lceil r/7 \rceil - $core of $G_2, \ldots, G_6$.
8.   **for** $2 \leq i \leq 6$ $G_i = (V_i, E_i) \leftarrow G_i[K]$.
9.   **end** SPLIT

We will show later (Lemma 4) that **whp** this algorithm terminates with $|K| \geq n - \epsilon n$ where ϵ is small. Note that SPLIT ensures that

- The final graphs G_i, $i = 2, \ldots, 6$ have the same vertex set K.
- Every $v \in K$ has degree at least $\lceil r/7 \rceil$ in each G_i, $1 \leq i \leq 6$.
- Every $v \in K$ has degree at most $2r/7$ in each G_i, $2 \leq i \leq 6$.
- If $v \in V \setminus K$ then $d_{G_1}(v) = d_G(v) = r$.

4.2 Phase 1.

Choose $z_1, z_2, \ldots, z_{2\kappa} \in V_2$ uniformly and randomly *without* replacement. Let Z denote the set $\{z_1, z_2, \ldots, z_{2\kappa}\}$. We are going to replace the problem of finding paths from a_i to b_i by that of finding paths from $\tilde{a}_i$ to $\tilde{b}_i$, where $\{\tilde{a}_1, \tilde{b}_1, \tilde{a}_2, \tilde{b}_2, \ldots, \tilde{a}_\kappa, \tilde{b}_\kappa\} = Z$. Let A denote the set $\{a_1, b_1, a_2, b_2, \ldots, a_\kappa, b_\kappa\}$.

We connect A to Z via edge-disjoint paths in the graph G_1 using network flow techniques. We construct a network as follows

- Each undirected edge of G_1 gets capacity 1.
- Each $v \in V$ becomes a source of capacity $|\{i : a_i = v \text{ or } b_i = v\}|$ and each member of Z becomes a sink of capacity 1.

Then we find a flow from A to Z that satisfies all demands. Since the maximum flow has integer values, it decomposes naturally into $|A|$ edge-disjoint paths (together perhaps with some cycles). If a path joins a_i to $z \in Z$, then we let $\tilde{a}_i = z$. Similarly, if a path joins b_i to $z \in Z$, then we let $\tilde{b}_i = z$.

Thus Phase 1 finds edge-disjoint paths $W_i^{(1)}$ from a_i to $\tilde{a}_i$ and $W_i^{(6)}$ from $\tilde{b}_i$ to b_i, $1 \leq i \leq \kappa$, where the vertices $\tilde{a}_1, \tilde{b}_1, \tilde{a}_2, \tilde{b}_2, \ldots, \tilde{a}_\kappa, \tilde{b}_\kappa \in V_2$ are chosen uniformly at random without replacement. On the other hand there may be some difficult conditioning involved in the pairing of $\tilde{a}_i$ with $\tilde{b}_i$, $1 \leq i \leq \kappa$. We deal with this in Phase 2(a).

4.3 Phase 2.

4.3.1 Algorithm GENPATHS.

We construct edge-disjoint paths connecting $\tilde{a}_i, \tilde{b}_i$ for $1 \leq i \leq \kappa$ in this phase using random walks. A *random walk* on an undirected graph (or multigraph) $G = (V, E)$ is a Markov chain $\{X_t\}$ on V associated with a particle that moves from vertex to vertex according to the following rule: The probability of a transition from vertex v, of degree d_v to a vertex w is $1/d_v$ if $\{v, w\} \in E$ and 0 otherwise. (For multigraphs, each edge out of a vertex is an equally likely exit; loops are counted as two exits.) Its stationary distribution, denoted by π or $\pi(G)$, is given by

$$\pi_v = \frac{d_v}{2|E|}.$$

A trajectory W of length τ is a sequence of vertices $[w_0, w_1, \ldots, w_\tau]$ such that $\{w_t, w_{t+1}\} \in E$ for $1 \leq t < \tau$. The Markov chain induces a probability distribution on trajectories in the usual way. We use $\mathbf{Pr}_G^{(\tau)}(a, b)$ to denote the probability that a random walk in G of length τ starting at a terminates at b.

There is some possible conditioning involved in the pairing of $\tilde{a}_i$ with $\tilde{b}_i$. To break this we first do random walks $W_i^{(2)}, W_i^{(5)}$ starting from $\tilde{a}_i, \tilde{b}_i, 1 \leq i \leq \kappa$ in graph Γ_2. Here $\Gamma_j = (K_j, E_j), j = 2, \ldots, 5$ denotes G_j after the deletion of some vertices and edges. We do these in the order $z_1, z_2, \ldots, z_{2\kappa}$. The walks are sufficiently long that their endpoints $\hat{a}_i, \hat{b}_i$ are essentially independent of $\tilde{a}_i$ and $\tilde{b}_i$. The edges of each such a walk are deleted before the next walk starts. This keeps the walks edge-disjoint. This constitutes Phase 2(a).

Then for $i = 1, \ldots, \kappa$ we do the following: We construct a random walk $W_i^{(4)}$ starting at $\hat{b}_i$ in Γ_4 of length τ_0, the endpoint is denoted by b_i^*. (This might seem unnecessary and we explain why we do it later.) Then in Step 20, we execute the subroutine WALK to generate a

set of random walks $\hat{W}_1, \ldots, \hat{W}_l$ starting at $\hat{a}_i$ in Γ_3. The last walk $\hat{W}_l$ will end at b_i^* and be accepted as $W_i^{(3)}$. Finally, we delete $\hat{W}_1, \ldots, \hat{W}_l$ from G_3 and delete $W_i^{(4)}$ from G_2.

It is important for the above analysis to ensure that our random walks take place on graphs with sufficiently high minimum degree. The obvious reason for this is that we do not wish to isolate any vertex that needs to be an endpoint of a walk. We also require that the walks take place on graphs which are strong expanders. For us, this requires a high minimum degree. In general we ensure that graphs supporting the random walks have minimum degree at least $r/10$. We manage this by the simple artifact of removing vertices of low degree until the graph has none left. (We could empty a graph this way, but we show later that this does not often happen.)

If the proposed start vertex v of a walk is of low degree then we try to *connect it back* to vertices of large degree (those in K_j) by a path in G_5. The terminal endpoint of this walk is denoted by v'. (Note that this paths is not of length τ_0, we expect it to be much shorter). We use a subroutine CONNECTBACK for this purpose. We do not expect to succeed all the time and our failures are kept in a set L for later consideration.

The walk from a_i to b_i is then the catenation of walks $W_i^{(t)}$, $t = 1, \ldots, 6$. These walks may each include a short walk W_{CB} at the beginning provided by CONNECTBACK.

For $z \in Z$ we use the notation $\phi(z) = i$ to indicate $z \in \{\tilde{a}_i, \check{b}_i\}$.

1. **Algorithm** GenPaths
2. **begin**
3.   **for** $i = 1$ **to** 2κ **do**
4. $K_2 \leftarrow \lceil r/10 \rceil$-core of G_2, $\Gamma_2 \leftarrow G_2[K_2]$.
5.   **if** $z_i \notin K_2$ **then**
6.   **Execute** ConnectBack$(K_2, z_i, z'_i, \phi(z_i), W_{CB})$ **fi**
7.   **if** $j = \phi(z_i) \notin L$ **then**
8.   Let $\theta = 2$ if $z_i = \tilde{a}_j$ and $\theta = 5$ if $z_i = \tilde{b}_j$. Construct
        random walk $W_j^{(\theta)}$ starting at z'_i in Γ_2. The terminal
        endpoint is $\hat{a}_j$ if $z_i = \tilde{a}_j$ and $\hat{b}_j$ otherwise.
9. $W_j^{(\theta)} \leftarrow (W_{CB}, W_j^{(\theta)}), G_2 \leftarrow G_2 \setminus E(W_j^{(\theta)})$
10.   **fi**
11.   **od**
12.   **for** $i = 1$ **to** κ **do**
13. $K_3 = K_4 \leftarrow \lceil r/10 \rceil$-core of G_3, G_4, $\Gamma_j \leftarrow G_j[K_j]$, $j = 3, 4$.
14.   **if** $\hat{a}_i \notin K_3$ **then**
15.   **Execute** ConnectBack $(K_3, \hat{a}_i, \hat{a}'_i, i, W_{CB}^{(3)})$ **fi**
16.   **if** $\hat{b}_i \notin K_4$ **then**
17.   **Execute** ConnectBack $(K_4, \hat{b}_i, \hat{b}'_i, i, W_{CB}^{(4)})$ **fi**
18.   **if** $i \notin L$ **then**
19.   Construct random walk $W_i^{(4)}$ starting at $\hat{b}'_i$ in Γ_4, the
        endpoint is denoted by b_i^*.
20.   **Execute** Walk$(\hat{a}'_i, \hat{b}'_i, b_i^*, \Gamma_3, \Gamma_4)$, $W_i^{(3)} \leftarrow \hat{W}_l$.
21. $G_3 \leftarrow G_3 \setminus E(\hat{W}_j), j = 1, \ldots, l$; $G_4 \leftarrow G_4 \setminus E(W_i^{(4)})$
22. $W_i^{(3)} \leftarrow (W_{CB}^{(3)}, W_i^{(3)})$, $W_i^{(4)} \leftarrow (W_{CB}^{(4)}, W_i^{(4)})$.
23.   **fi**
24.   **od**
25. **end** GenPaths

4.3.2 Subroutine ConnectBack.

The purpose of ConnectBack is to connect a vertex x to vertices of large degree. All walks are done on vertices of large degree and in Steps 3–6 we check that the start point x has large degree in Γ_5. (Γ_5 is initially G_5 but loses vertices and edges as the algorithm proceeds). If not, we put it into L for later. We do an *edge avoiding* random walk W_{CB} from x until a vertex x' of K is reached. Later we replace x by x'. By edge avoiding, we mean that one is not allowed to use the same edge more than once. We do this for technical reasons, but it does not seem unreasonable anyway. This is done in Steps 7–10.

1.   **subroutine** CONNECTBACK(K, x, x', i, W_{CB})

2.   **begin**

3. $\Gamma_5 \leftarrow \lceil r/10 \rceil$-core of G_5

4.   **if** $x \notin \Gamma_5$ **then**

5. $L \leftarrow L \cup \{i\}$

6.   **exit**

7.   **fi**

8.   Construct an edge avoiding random walk W_{CB} starting at x in Γ_5 until
                    W_{CB} visits some vertex x' in K

9.   (This is a trivial walk if $x \in K$).

10. $G_5 \leftarrow G_5 \setminus W_{CB}$

11.   **end** CONNECTBACK

4.3.3 Subroutine WALK.

To connect $\hat{a}_i$ and $\hat{b}_i$ after Phase 2(a) we first connect them (if necessary) to the $\lceil r/10 \rceil$-core
of G_3, G_4. We then do a random walk $W_i^{(4)}$ in Γ_4 starting at $\hat{b}_i'$ and ending at b_i^*. We then try
to join $\hat{a}_i'$ to b_i^*.

   To join $\hat{a}_i'$ and b_i^*, the reader might expect us to choose a random walk from those with
endpoints $\hat{a}_i', b_i^*$. The main problem with this is that the distribution of b_i^* may be significantly
different from the steady state distribution of a walk from $\hat{a}_i'$ in Γ_3. If we choose a walk in
this manner then deleting it will condition the graph in a way which is complex to analyze,
especially as we have to repeat the procedure κ times.

   We overcome this by choosing a set of random walks and use rejection sampling to make
the final walk have the correct distribution. There is still the complication that the b_i^*
are chosen before we do the walks. This leads to the subroutine WALK described next.
WALK$(\hat{a}_i', \hat{b}_i', b_i^*, \Gamma_3, \Gamma_4)$ generates a series of random walks of length τ_0 in Γ_3 starting from
$\hat{a}_i'$. The last walk generated ends at b_i^* which has the distribution

$$\hat{p}_v = P_{\Gamma_4}^{(\tau_0)}(\hat{b}_i', v).$$

   [There is the question as to why we do not try to construct a walk from $\hat{a}_i$ to $\hat{b}_i$. We need
in WALK to generate an endpoint from the same distribution as the target. This will not be
so easy if $\hat{b}_i$ has dropped out of the $\lceil r/10 \rceil$-core.]

   The somewhat strange method used to generate these walks will be further explained in
Section 8.

1.   **subroutine** $\textsc{Walk}(\hat{a}'_i, \hat{b}'_i, b^*_i, \Gamma_3, \Gamma_4)$

2.   **begin**

3. $p_v \leftarrow P^{(\tau_0)}_{\Gamma_3}(\hat{a}'_i, v)$ for $v \in K_3$

4. $\hat{p}_v \leftarrow P^{(\tau_0)}_{\Gamma_4}(\hat{b}'_i, v)$ for $v \in K_4$ (the distribution of b^*_i.)

5. $p_{\min} \leftarrow \min\{p_v : v \in K_3\}$

6. $\hat{p}_{\max} \leftarrow \max\{\hat{p}_v : v \in K_4\}$

7.   Choose l from the geometric distribution with probability of success
   $$s = p_{\min}/\hat{p}_{\max}$$

8.   **for** k **from** 1 **to** $l-1$ **do**

9.   Choose x_k according to $\mathbf{Pr}(x_k = v) = (p_v - \hat{p}_v p_{\min}/\hat{p}_{\max})/(1-s)$

10.   **od**

11. $x_l \leftarrow b^*_i$

12.   **for** k **from** 1 **to** l **do**

13.   Pick a walk $\hat{W}_k$ of length τ_0 in Γ_3 according to the distribution
    on trajectories, conditioned on start point $= \hat{a}'_i$ and
    end point $= x_k$

14.   **od**

15.   **output** $\hat{W}_1, \hat{W}_2, \ldots, \hat{W}_l$

16.   **end** $\textsc{Walk}$

The distributions p_v and $\hat{p}_v$ can be computed in $O(nm\tau_0)$ time by computing powers of the transition matrix, after which a random walk can be found in $O(n\tau_0)$ time. (For details see [10].) The analysis will show that in the range of interest, **whp**, s is bounded away from zero by a constant, hence the expected total running time of $\textsc{Walk}$ is $O(nm\tau_0)$.

4.4   Phase 3.

There is still the set L of pairs which have not been connected by paths. We will show later that **whp** $|L|$ is at most $n/(\log n)^{C_0}$. As such, these pairs can be dealt with by the algorithm of [10], using graph G_6.

5   Configuration model

5.1   Model.

We first show (Lemma 2 below) that the graphs G_i, $i = 2, \ldots, 6$ are random given their degree sequences.

    The simplest model for graphs with a fixed degree sequence is the *configuration model* of Bollobás [4] which is a probabilistic interpretation of the counting formula of Bender and Canfield [2]. Let $\mathbf{d} = \{d_1, d_2, \ldots, d_\nu\}$ denote a degree sequence, *block $D_i = \{1, \ldots, d_i\} \times \{i\}$* for $1 \leq i \leq \nu$ and $D = \cup_{i=1}^{\nu} D_i$ be the set of *points*. Let $\Omega = \Omega(D)$ be the set of partitions of D into pairs. Such partitions are called *configurations*. If $F \in \Omega$ then the multigraph $M(F)$ is defined as follows: $V(M) = [\nu]$ and there is an edge $\{i, j\}$ for every pair in F of the form $\{(x, i), (y, j)\}$ (for some x and y). Assume F is chosen randomly from Ω. The salient properties of $M = M(F)$ (see [4]) are:

Lemma 1 (a) *If M is simple, then it is equally likely to be any simple graph with degree sequence $\mathbf{d}$.*

(b) $\mathbf{Pr}(M$ is simple$)$ $= \exp\{-O(\mu^2/\nu^2)\}$ *where $\mu = |D|/2$ is the number of edges in M – hence $2\mu/\nu$ is the average degree of M.*

It therefore suffices to replace $G_{n,r}$ by a random configuration multigraph $G = M(F)$ where $\mathbf{d} = (r, r, \ldots, r)$. The corresponding random partition is denoted by F_0. Any property that holds **whp** in this model, holds **whp** in $G_{n,r}$ since $\exp\{-O(\mu^2/\nu^2)\} = e^{-O(1)}$. In what follows the graphs $G_2, \ldots, G_6$ are multigraphs associated with configurations $F_2, \ldots, F_6$.

5.2 Random walks on configurations.

To analyse random walks on the graphs G_i, we need to be sure they have the correct distribution.

Lemma 2 *For each $i = 2, \ldots, 6$ the configuration F_i produced by algorithm SPLIT is random given its degree sequence i.e. conditional on having degree sequence $\mathbf{d}$, F_i is equally likely to be any pairing of the corresponding $\sum d_i$ points.*

Proof: Let F' be obtained from F by including each pair of F with probability $1/6$. We first show that F' is a random configuration given its degree sequence. Notice that

$$\mathbf{Pr}(F' \mid \mathbf{d}) = \frac{\mathbf{Pr}(F')}{\mathbf{Pr}(\mathbf{d})} = \frac{\sum_F \mathbf{Pr}(F' \mid F)}{|\Omega|\,\mathbf{Pr}(\mathbf{d})}$$

and $\mathbf{Pr}(F' \mid F) = (1/6)^{(\sum d_i)/2}(5/6)^{(rn - \sum d_i)/2}$ if $F' \subset F$ and 0 otherwise. Suppose F_1 and F_2 are two configurations with the same degree sequence $\mathbf{d}$, and let $\mathcal{F}_1 = \{F \mid F_1 \subset F\}$ and $\mathcal{F}_2 = \{F \mid F_2 \subset F\}$. It suffices to show that $|\mathcal{F}_1| = |\mathcal{F}_2|$. Let E_0 be the set of pairs $F_1 \oplus F_2 = (F_1 \setminus F_2) \cup (F_2 \setminus F_1)$. Define mapping $f : \mathcal{F}_1 \to \mathcal{F}_2$ as follows: Given $F \in \mathcal{F}_1$, let $f(F) = [F \setminus (E_0 \cap F_1)] \cup (E_0 \cap F_2)$. It's easy to see that $f(F) \in \mathcal{F}_2$ and $f^{-1} : \mathcal{F}_2 \to \mathcal{F}_1$ is well-defined as $f^{-1}(F) = [F \setminus (E_0 \cap F_2)] \cup (E_0 \cap F_1)$. Hence f is a bijection and $|\mathcal{F}_1| = |\mathcal{F}_2|$.

Thus (in an obvious notation) $F_1', \ldots, F_6'$ are all random configurations. Deleting a set D_j and all t pairs containing a member of D_j from F_i' preserves randomness since every derived configuration arises from the same number of possible configurations F_i'. So the final configurations $F_2, \ldots, F_6$ are random, given their degree sequences. $\quad\square$

We need to show that removing the edges of a random walk does not condition the pairings of the remaining points. We now consider the construction of a random configuration F. It is useful to think of F as being constructed sequentially.

1. **Algorithm** CONSTRUCT
2. **begin**
3. $\qquad F_0 \leftarrow \emptyset;\ R_0 \leftarrow W$
4. $\qquad$ **for** $t = 1$ **to** $|W|/2$ **do**
5. $\qquad\qquad$ Choose $u_t \in R_{t-1}$ *arbitrarily*
6. $\qquad\qquad$ Choose v_t *randomly* from $R_{t-1} \setminus \{u_t\}$
7. $\qquad\qquad F_t \leftarrow F_{t-1} \cup \{\{u_t, v_t\}\};\ R_t \leftarrow R_{t-1} \setminus \{u_t, v_t\}$
8. $\qquad$ **od**
9. $\quad$ **output** F
10. $\quad$ **end** CONSTRUCT

It is important to observe that for any $t > 0$, $F \setminus F_t$ is a random member of $\Omega(R_t)$.

An important consequence of the above observation is that if we start with $M = M(F)$, then the multigraph obtained by removing from M the edges of a random walk W remains unconditioned. Indeed, we may imagine CONSTRUCT as performed in parallel with our walk W. Suppose our walk makes a transition from a vertex x and the current value of R_t in CONSTRUCT is R. The transition from x is equivalent to choosing a random member $u = u_t$ of D_x. If $u \in R$, then we perform one step of CONSTRUCT and pair u with a point $v = v_t \in R \setminus \{u\}$. If $v \in D_y$ for some y, then the walk makes a transition from x to y. If $u \notin R$ then v is the point already paired with u. Thus, since $F \setminus F_t$ is random, we see that removing from M the edges of a random walk results in a multigraph from a *random* configuration. The same argument can be applied to the edge avoiding walk of CONNECTBACK. In the latter case we do not need to consider the case $u \notin R_t$.

In summary, we can conduct our analysis of Phase 3 always assuming that the graphs G_2 – G_5 are random configuration multigraphs, given their degree sequences.

6   Analysis of Phase 1

We begin with some discussion of G:

Lemma 3 Whp *every subgraph H of G which is induced by a set of vertices S, $|S| \leq n_0 = 27n/(8e^6 r^3)$, has at most $3|S|/2$ edges.*

Proof: Let $D_S = \bigcup_{i \in S} D_i = \{w_1, w_2, \ldots, w_{r|S|}\}$. Let us randomly pair up $w_1, w_2, \ldots$. At any stage, if w_i has not been paired with any of $w_1, w_2, \ldots, w_{i-1}$, and given the pairings of these points, the probability that w_i is paired with something in D_S is at most $|S|/n$. Hence

$$
\begin{aligned}
\mathbf{Pr}(\exists S : \ |S| = s \leq n_0, \ S \text{ contains } 3s/2 \text{ edges}) &\leq \sum_{s=4}^{n_0} \binom{n}{s} \binom{rs}{3s/2} \left(\frac{s}{n}\right)^{3s/2} \\
&\leq \sum_{s=4}^{n_0} \left(\frac{ne}{s}\right)^s \left(\frac{2ers}{3n}\right)^{3s/2} \\
&= \sum_{s=4}^{n_0} \left(\frac{8e^5 r^3 s}{27n}\right)^{s/2} \\
&= O(n^{-2}).
\end{aligned}
$$

$\square$

The next lemma concerns the size of K at the end of SPLIT.

Lemma 4 *If K is as defined in algorithm SPLIT then* **whp**

$$
|K| \geq (1 - e^{-r/800})n.
$$

Proof: Fix $1 \leq i \leq 6$ and consider G'_i immediately after Step 4. If $v \in V$ then $d_{G'_i}(v)$ has distribution $\text{Bin}(r, 1/6)$ and hence at this point (1) implies

$$
\mathbf{Pr}(d_{G'_i}(v) \leq \lfloor r/7 \rfloor) \leq \epsilon_1 = e^{-r/600}.
$$

So if $V_i = \{v \in V : \ d_{G'_i}(v) > \lfloor r/7 \rfloor\}$ then $\mathbf{E}(|V_i|) \geq (1 - \epsilon_1)n$. Now the random variable $|V_i|$ is determined by the outcome of $rn/2$ random choices. Also, changing one such choice results

in changing $|V_i|$ by at most 2. Applying the Azuma-Hoeffding martingale tail inequality ([5], [21]) we obtain

$$\mathbf{Pr}(|V_i| \leq \mathbf{E}(|V_i|) - t) \leq e^{-t^2/(rn)}.$$

Putting $t = \epsilon_1 n$ we see that **whp**

$$|V_i| \geq (1 - 2\epsilon_1)n, \qquad 1 \leq i \leq 6. \tag{4}$$

Now consider K at the end of Step 5. Let $W_0 = V \setminus V_1$ and define a sequence $w_1, w_2, \ldots, w_k$ as follows: Suppose $W_j = W_0 \cup \{w_1, w_2, \ldots, w_j\}$ and w_{j+1} is an arbitrary vertex not in W_j but which has at least 2 neighbours in W_j. The sequence stops at index k when every vertex not in W_k has at most one neighbour in W_k. Now every vertex not in W_k has $d_{G'_i}(v) > \lfloor r/7 \rfloor$ and hence has degree at least $\lfloor r/7 \rfloor$ in $V \setminus W_k$. Thus $|K| \geq n - |W_k|$. We now show that **whp** $k < 6\epsilon_1 n$. If not then

- $W_{6\epsilon_1 n}$ has at most $8\epsilon_1 n$ vertices.

- $W_{6\epsilon_1 n}$ contains at least $12\epsilon_1 n$ edges.

For sufficiently large r, this contradicts the conclusion of Lemma 3. Thus **whp**, at the end of Step 5,

$$|K| \geq (1 - 8\epsilon_1)n. \tag{5}$$

Now consider the execution of Step 7. One can see that the value for K at the termination of this step is independent of the order in which vertices are deleted by the k-core algorithm. So assume that we start by deleting $\bigcup_{i=2}^{6}(V \setminus V_i)$. (4) and (5) imply that **whp** after this $|K| \geq (1 - 18\epsilon_1)n$. We then argue, similarly to above, that no more than $72\epsilon_1 n$ more vertices will be deleted by the k-core algorithm. The lemma follows for sufficiently large r. $\quad \square$

7 Analysis of Phase 2

In this section we show that if our input graph $G = (V, E)$ is $G_{n,r}$ then **whp**, after we run SPLIT, we can find in G_1 edge-disjoint paths from a_i to $\tilde{a}_i$, and b_i to $\tilde{b}_i$, for $1 \leq i \leq \kappa$, for *any* choice of $a_1, \ldots, b_\kappa$ consistent with the premises of Theorem 1, and every choice for $\tilde{a}_1, \ldots, \tilde{b}_\kappa$.

Let A and Z be as defined in Section 4.2. For $S \subseteq V$, let

$$\alpha(S) = |\{i : a_i \text{ or } b_i \in S\}| \text{ and } \xi(S) = |S \cap Z|.$$

For sets $S, T \subseteq V$, let $e_{G_1}(S, T)$ denote the number of edges of G_1 with an endpoint in S and the other endpoint in T. It suffices to prove that

$$e_{G_1}(S, \bar{S}) \geq \xi(\bar{S}) - \alpha(\bar{S}), \qquad \forall S \subseteq V. \tag{6}$$

We can then apply a theorem of Gale [14] (see Bondy and Murty [7], Theorem 11.8) to deduce the existence of the required flow in G_1 for the successful run of Phase 2.

Lemma 5 Whp $e_{G_1}(S, \bar{S}) \geq r|S|/288$ *for all* $S \subseteq V$, $|S| \leq n/2$.

Proof: Assume first that $|S| \leq n_0$. From Lemma 3, we obtain

$$e_G(S, \bar{S}) \geq (\delta(G_1) - 3)|S| \geq (\lceil r/7 \rceil - 3)|S|.$$

So now assume $|S| > n_0$ and first consider $e_G(S, \bar{S})$. We claim that in the configuration model this stochastically dominates $\text{Bin}(r|S|/4, 1/3)$. Apply algorithm CONSTRUCT, always choosing $u_t \in D_v$, $v \in S$ whenever possible. This is definitely possible in the first $r|S|/4$ choices. For each such choice, the probability that $v_t \in D_w$, $w \in \bar{S}$ is at least $1/3$, since $|R_{t-1} \cap S| \leq r|S| \leq rn/2$ and $|R_{t-1} \cap \bar{S}| \geq rn/4$. But by (1) with $\epsilon = 1/2$,

$$\mathbf{Pr}\left(\text{Bin}\left(\frac{r|S|}{4}, \frac{1}{3}\right) \leq \frac{r|S|}{24}\right) \leq e^{-r|S|/96}.$$

Thus

$$\mathbf{Pr}\left(\exists S : \ n_0 \leq |S| \leq n/2 \text{ and } e_G(S, \bar{S}) \leq \frac{r|S|}{24}\right) \leq \sum_{s=n_0}^{n/2} \binom{n}{s} e^{-rs/96}$$

$$\leq \sum_{s=n_0}^{n/2} \left(\frac{ne}{s} e^{-r/96}\right)^s$$

$$= o(1)$$

for r sufficiently large.

Assume then that $e_G(S, \bar{S}) \geq r|S|/24$ for $n_0 \leq |S| \leq n/2$. Given this, $e_{G_1}(S, \bar{S})$ stochastically dominates $\text{Bin}(r|S|/24, 1/6)$. Applying (1) again with $\epsilon = 1/2$,

$$\mathbf{Pr}\left(\exists S : \ n_0 \leq |S| \leq n/2 \text{ and } e_G(S, \bar{S}) \leq \frac{r|S|}{288}\right) \leq \sum_{s=n_0}^{n/2} \binom{n}{s} e^{-rs/1152}$$

$$\leq \sum_{s=n_0}^{n/2} \left(\frac{ne}{s} e^{-r/1152}\right)^s$$

$$= o(1)$$

for r sufficiently large. $\quad\square$

Now

$$\alpha(\bar{S}) = 2\kappa - \alpha(S) \geq 2\kappa - \beta r|S|$$

and so

$$\xi(\bar{S}) - \alpha(\bar{S}) \leq |Z \cap \bar{S}| - 2\kappa + \beta r|S| \leq \beta r|S|.$$

Thus Lemma 5 verifies (6) for $|S| \leq n/2$ provided we choose $\beta < 1/288$. For $|S| > n/2$ we use

$$e_{G_1}(S, \bar{S}) = e_{G_1}(\bar{S}, S) \geq r|\bar{S}|/288 \geq |Z \cap \bar{S}| \geq |Z \cap \bar{S}| - 2\kappa + \alpha(S) = \xi(\bar{S}) - \alpha(\bar{S})$$

and so Phase 2 succeeds **whp**.

8   Analysis of Phase 3

8.1   Construction of $W_i^{(2)}, W_i^{(5)}$

For $v \in K$, let $d_{j,v}$ denote its degree in $\Gamma_j, j = 2, 3, 4$ and let $\pi_{j,v} = d_{j,v} / \sum_{w \in K} d_{j,w}$ be the steady state distribution of a random walk on Γ_j. These values change as the algorithm progresses, but immediately prior to any random walk in a Γ_j we ensure that

$$\frac{r}{10} \leq d_{j,v} \leq \frac{2r}{7}. \tag{7}$$

Hence $v \in K$ implies

$$\frac{1}{3|K|} \leq \pi_{j,v} \leq \frac{3}{|K|}. \tag{8}$$

Let $\Gamma^{(i)}$ denote Γ_2 immediately prior to constructing the random walk starting at z_i and let $K^{(i)}$ be its vertex set, the current value of K. Let $P^{(i)}$ denote the transition probability matrix of a random walk on $\Gamma^{(i)}$. Let $\lambda^{(i)}$ be the second largest eigenvalue of $P^{(i)}$. It is well known that the second eigenvalue determines the rate of convergence of a Markov chain to its steady state. An explicit form of this result was obtained by Jerrum and Sinclair [25]: if $P_{\Gamma^{(i)}}^{(t)}(u,v)$ denotes the probability that a random walk of length t in $\Gamma^{(i)}$ which starts at u will end at v, then

$$|P_{\Gamma^{(i)}}^{(t)}(u,v) - \pi_v^{(i)}| \leq \left(\lambda^{(i)}\right)^t \sqrt{\frac{\pi_v^{(i)}}{\pi_u^{(i)}}}. \tag{9}$$

We will argue later that **whp** throughout the algorithm

$$|K| \geq (1 - \epsilon_K)n, \tag{10}$$

where

$$\epsilon_K = 8e^{-r/250000}.$$

We will argue that as long as K is this large, vertices are being removed at a sufficiently slow rate so that the overall loss from K is less than $\epsilon_K n$ **whp**. Thus the specific assumption at each stage is that (10) holds at the present time. We do not in this way "condition on the future".

We argued in Section 5.2 that $\Gamma^{(i)}$ is random, given its degree sequence. According to Theorem 2 of [8], (which is adapted from [16]) the second largest eigenvalue $\lambda^{(i)}$ is with probability $1 - O(|K|^{-a}) = 1 - O(n^{-a})$ bounded by $\gamma/\sqrt{r}$, where γ depends only on a and a bound on the ratio of the maximum and the minimum degrees of $\Gamma^{(i)}$, which we can take to be 3 by (7). Taking $a = 4$ we have that γ is a constant independent of r. Using this in (9) we obtain that conditional on an event of probability $1 - O(n^{-4})$

$$|P_{\Gamma^{(i)}}^{(t)}(u,v) - \pi_v^{(i)}| \leq \frac{3\gamma^t}{r^{t/2}}. \tag{11}$$

So from (3) we see that for large r we have

$$|P_{\Gamma^{(i)}}^{(\tau_0)}(u,v) - \pi_v^{(i)}| \leq n^{-4}. \tag{12}$$

Let $\mathcal{E}^{(i)}$ be the intersection of the high probability events described in Lemmas 3, 4 and 5 plus (10), (12) for $\Gamma^{(1)}, \ldots, \Gamma^{(i)}$. We prove by induction on i that

$$\mathbf{Pr}(\mathcal{E}^{(i+1)} \mid \mathcal{E}^{(i)}) \geq 1 - O(n^{-3}). \tag{13}$$

This is clearly true for $i = 0$ and all we need to show is that (10) holds with the correct probability.

Fix $1 \le i \le \kappa$, let $s = 100 \ln n$ and $\hat{V}_2$ be the set of vertices in G_2 which are incident with at least $r/30$ edges of $W_1, \ldots, W_i$. For $S \subseteq K^{(0)}$ of size s, let $\mathcal{B}_S$ be the event that $S \subseteq \hat{V}_2$. We first show that for all such S,

$$\mathbf{Pr}(\mathcal{B}_S \mid \mathcal{E}^{(i)}) \le 2e^{-rs/250000}. \tag{14}$$

Let $S = \{v_1, \ldots, v_s\} \subset K^{(0)}$. Let $\mathcal{P} = \mathcal{P}(S)$ be the event that there exist $w \in K^{(1)} \setminus S$ and $t = r/2000$ distinct paths $P_1, P_2, \ldots, P_t$ from w to S of length $\le k_0 = (\log_r n)/10$. If such paths exist then the graph induced by $w, S, P_1, P_2, \ldots, P_t$ contains $s + u$ vertices and $\ge s + u + t - 1$ edges for some $1 \le u \le (k_0 - 1)t + 1$. This is because the addition of each new path to S, w adds at least one more new edge than vertex. Thus,

$$\begin{aligned}
\mathbf{Pr}(\mathcal{P}) &\le \sum_{u=1}^{(k_0-1)t+1} \binom{n}{s+u} \binom{2r(s+u)/7}{s+u+t-1} \left(\frac{3(s+u)}{n}\right)^{s+u+t-1} \tag{15} \\
&\le \sum_{u=1}^{(k_0-1)t+1} \left(\frac{ne}{s+u}\right)^{s+u} \left(\frac{r(s+u)e}{n}\right)^{s+u+t-1} \\
&\le \sum_{u=1}^{(k_0-1)t+1} \frac{(re^2)^{s+u+t}(s+u)^{t-1}}{n^{t-1}} \\
&\le \frac{2(re^2)^{s+k_0t+1}(s+k_0t)^{t-1}}{n^{t-1}} \\
&\le n^{-r/2500} \\
&= e^{-rs/250000} \tag{16}
\end{aligned}$$

Explanation of (15): working within the configuration model we choose the $s + u$ vertices in at most $\binom{n}{s+u}$ ways. The corresponding blocks contain at most $2r(s+u)/7$ points X. We then choose the set of lowest indices I of the $s + u + t - 1$ pairs in at most $\binom{2r(s+u)/7}{s+u+t-1}$ ways. The last factor is an upper bound on the probability that each point in I chooses a partner in X.

Assume $\mathcal{P}$ does not occur. Let W_i denote the walk $W_i^{(2)}$. Let $Z_{i,S}$ denote the number of edges incident with vertices in S that are used by W_i and let $N_{i,S}$ be the number of visits to S by W_i. Consider ρ_S, the probability that a random walk W from $u \in S$ returns to S within τ_0 steps.

Lemma 6 *If $\mathcal{P}$ does not occur then $\rho_S \le 1/20$.*

Proof: Let $dist(v, S)$ denote the length of a shortest path in $\Gamma^{(i)}$ from v to S. Define $S_j = \{v \in K^{(i)} \setminus S : dist(v, S) = j\}$ for $j = 1, \ldots, k_0 - 1$ and $S_{k_0} = \{v \in K^{(i)} \setminus S : dist(v, S) \ge k_0\}$.

We define a random walk $X_0, X_1, \ldots$, where $X_0 = 0$ and $X_t = j$ iff W is at a vertex $v \in S_j$ after t steps ($S_0 = S$). Then given $\bar{\mathcal{P}}$ and that the minimum degree in $\Gamma^{(i)}$ is $r/10$, for a random walk W,

$$\mathbf{Pr}(X_t = j + 1 \mid X_t = j) \ge 199/200, \qquad j = 1, \ldots, k_0 - 1 \tag{17}$$
$$\mathbf{Pr}(X_{t+1} = k_0 \mid X_t = k_0) \ge 199/200. \tag{18}$$

Now define another random walk $\tilde{X}_0 = 0, \tilde{X}_1, \ldots,$ where

$$\mathbf{Pr}(\tilde{X}_{t+1} = \tilde{X}_t + \delta) = \begin{cases} \dfrac{199}{200} & \delta = 1 \\ \dfrac{1}{200} & \delta = -1 \ \tilde{X}_t > 0, \text{ or} \\ \dfrac{1}{200} & \delta = 0 \quad \tilde{X}_t = 0 \end{cases}$$

Now for $u > 0$ we have

$$\mathbf{Pr}(\tilde{X}_u = 0) \le \binom{u}{\lceil u/2 \rceil} \left(\frac{1}{200}\right)^{\lceil u/2 \rceil} \le 50^{-\lceil u/2 \rceil}$$

since if $\tilde{X}_u = 0$ at least $\lceil u/2 \rceil$ of the first u moves must be to the left i.e. $\delta \le 0$. So,

$$\mathbf{Pr}(\exists\, u > 0 : X_u = 0) \le \sum_{u=1}^{\infty} 50^{-\lceil u/2 \rceil} \le 2 \sum_{u=1}^{\infty} 50^{-u} = \frac{2}{49}.$$

It follows from (17) and (18) that $\tilde{X}_j \le X_j$ in distribution up to the first time u_0 that $X_u = k_0$. So
$$\mathbf{Pr}(\exists\, 0 < u \le u_0 : \ X_u = 0) \le \frac{2}{49}. \tag{19}$$

Next observe that for any $u > 0$

$$\mathbf{Pr}\left(X_{u+\frac{k_0}{10}} < \frac{k_0}{2} \mid X_u \ge \frac{k_0}{2}\right) \le 2^{k_0/10} 200^{-k_0/20} = 50^{-k_0/20}$$

since in order that $X_{u+\frac{k_0}{10}} < \dfrac{k_0}{2}$ there must be at least $\dfrac{k_0}{20}$ moves to the left.
   Putting $u_i = u_0 + i\frac{k_0}{10}$ we see that

$$\mathbf{Pr}\left(\exists\, i,\, 1 \le i \le 10 : \ X_{u_i} < \frac{k_0}{2}\right) \le 10 \times 50^{-k_0/20}.$$

Then observe that if $X_u = 0$ for some $u_0 < u \le \tau_0$ then we must have $u_i < \dfrac{k_0}{2}$ for some $1 \le i \le 10$. This together with (19) completes the proof of the lemma. $\quad\square$

   Let $q_k = \mathbf{Pr}(N_{i,S} = k \mid \mathcal{E}^{(i)})$ for $k \ge 1$. We claim that given $W_1, W_2, \ldots, W_{i-1}$,

$$q_k < \left(\frac{1000 s \log_r n}{n}\right) 20^{-(k-1)} \tag{20}$$

   To prove (20) for $k = 1$, fix $\Gamma^{(i)}$ and let $h_v(t)$ be the probability that the walk is at v after t steps. For $t = 0$, we have two cases.

Case 1: Subroutine CONNECTBACK is not executed. In this case W_i starts from z_i which is a vertex chosen uniformly at random from $V_2 \setminus \{z_1, z_2, \ldots, z_{i-1}\}$. Notice that the execution of Phase 2 does not depend on the paths created in Phase 1. Therefore we can analyse Phase 2 without reference to Phase 1 and it is legitimate to consider z_i as chosen randomly in this way, independent of the state of $\Gamma^{(i)}$. So

$$h_v(0) \le \frac{2}{|V_2|} \le 10\pi_v^{(i)}. \tag{21}$$

Case 2: Subroutine CONNECTBACK is executed. In this case W_i starts from z_I' which is the end point of a random walk in Γ_3. The random walks in Γ_3 only use unexposed edges. The endpoint is chosen randomly from $K^{(i)}$ with probability at most $2\pi_v^{(i)}$, see Section 8.3.

So in either case $h_v(0) \leq 10\pi_v^{(i)}$. We can now inductively show that $h_v(t) \leq 10\pi_v^{(i)}$ for all $t \geq 0$. Using the stationarity equations, we have

$$h_v(t+1) = \sum_{w \in N_{\Gamma^{(i)}}(v)} \frac{h_w(t)}{d_w^{(i)}} \leq 10 \sum_{w \in N_{\Gamma^{(i)}}(v)} \frac{\pi_w^{(i)}}{d_w^{(i)}} = 10\pi_v^{(i)} \leq \frac{100}{n}. \tag{22}$$

The last inequality follows from (8) and (10). Hence, if $\hat{q}_1$ denotes the probability that $N_{I,S} \geq 1$ then

$$q_1 \leq \hat{q}_1 \leq \sum_{v \in S} \sum_{t=0}^{\tau_0} h_v(t) \leq \frac{1000 s \log_r n}{n}.$$

For $k \geq 2$, (20) follows from

$$q_k \leq \rho_S^{k-1} \hat{q}_1.$$

We prove next that **whp** the size of $\hat{V}_2$ is small. First (20) implies that

$$\mathbf{E}(e^{N_{i,S}} \mid \mathcal{E}^{(i)}, \bar{\mathcal{P}}) \leq 1 + \sum_{k \geq 1} e^k \left(\frac{1000 s \log_r n}{n} \right) \rho_S^{k-1} \tag{23}$$

$$\leq 1 + \frac{C_2 s \log_r n}{n} \tag{24}$$

Clearly $Z_{i,S} \leq 2N_{i,S}$. Thus for any $t > 0$,

$$\mathbf{Pr}\left(\sum_{j=1}^{i} Z_{j,S} \geq st \;\middle|\; \mathcal{E}^{(i)}, \bar{\mathcal{P}} \right) \tag{25}$$

$$\leq e^{-st/2} \mathbf{E}\left(\exp\left\{ \sum_{j=1}^{i} N_{j,S} \right\} \;\middle|\; \mathcal{E}^{(i)}, \bar{\mathcal{P}} \right)$$

$$\leq e^{-st/2} \left(1 + \frac{C_2 s \log_r n}{n} \right) \mathbf{E}\left(\exp\left\{ \sum_{j=1}^{i-1} N_{j,S} \right\} \;\middle|\; \mathcal{E}^{(i)}, \bar{\mathcal{P}} \right)$$

$$\leq e^{-st/2} \exp\left\{ \frac{C_2 s \log_r n}{n} \right\} \mathbf{E}\left(\exp\left\{ \sum_{j=1}^{i-1} N_{j,S} \right\} \;\middle|\; \mathcal{E}^{(i-1)}, \bar{\mathcal{P}} \right) \frac{\mathbf{Pr}(\mathcal{E}^{(i-1)}, \bar{\mathcal{P}})}{\mathbf{Pr}(\mathcal{E}^{(i)}, \bar{\mathcal{P}})}$$

$$\leq \exp\left\{ -\frac{st}{2} + i\frac{C_2 s \log_r n}{n} \right\} \frac{1}{\mathbf{Pr}(\mathcal{E}^{(i)}, \bar{\mathcal{P}})}$$

$$\leq 2\exp\left\{ -\frac{st}{2} + 2\kappa \frac{C_2 s \log_r n}{n} \right\}$$

$$\leq 2\exp\{s(-t/2 + 2\alpha r C_2)\} \tag{26}$$

Take $t = r/30$ and $\alpha \leq (240C_2)^{-1}$ so that $-t + 2\alpha C_2 r < -r/120$. Therefore

$$\mathbf{Pr}(S \subseteq \hat{V}_2 \mid \mathcal{E}^{(i)}, \bar{\mathcal{P}}) \leq 2e^{-rs/120}.$$

Thus,

$$\mathbf{Pr}(S \subseteq \hat{V}_2 \mid \mathcal{E}^{(i)}) \leq \mathbf{Pr}(\mathcal{P} \mid \mathcal{E}^{(i)}) + \mathbf{Pr}(S \subseteq \hat{V}_2 \mid \mathcal{E}^{(i)}, \bar{\mathcal{P}}) \leq 2e^{-rs/250000}$$

proving (14).

So if $\epsilon_2 = e^{-r/250000}$ then it follows easily that **whp** $|\hat{V}_2| \leq 3\epsilon_2 n$. Indeed

$$\mathbf{Pr}(|\hat{V}_2| \geq 3\epsilon_2 n \mid \mathcal{E}^{(i)}) = \mathbf{Pr}\left(\left(\binom{|\hat{V}_2|}{s}\right) \geq \binom{3\epsilon_2 n}{s} \,\middle|\, \mathcal{E}^{(i)}\right)$$

$$\leq \frac{\mathbf{E}\left(\binom{|\hat{V}_2|}{s} \,\middle|\, \mathcal{E}^{(i)}\right)}{\binom{3\epsilon_2 n}{s}} \leq \frac{\binom{n}{s} 2\epsilon_2^s}{\binom{3\epsilon_2 n}{s}}$$

$$\leq 2(e/3)^s = O(n^{-9}). \tag{27}$$

Suppose that besides $\hat{V}_2$, we also delete $v_1, v_2, \ldots, v_\ell$ from K_2 in Step 4 of GENPATHS. We claim that $\ell \leq \epsilon_2 n$. To see this note that v_i must have at least $\lceil r/7 \rceil - r/30 - r/10 \geq r/200$ neighbours in $\hat{V}_2 \cup \{v_1, \ldots, v_{i-1}\}$. Thus $\hat{V}_2 \cup \{v_1, \ldots, v_{i-1}\}$ has at least $\ell r/200$ edges. This contradicts Lemma 3 for $\ell > \epsilon_2 n$ and r sufficiently large, which proves that $|K| \geq n(1 - 4\epsilon_2)$ **whp** for the duration of Phase 2(a).

8.2 Construction of $W_i^{(4)}$ and Analysis of WALK

The arguments we gave above apply to constructing and deleting $W_i^{(4)}$ in Γ_4 and we will focus on discussing the subroutine WALK. Consider a modification of WALK defined as follows:

1. **subroutine** WALK1$(\hat{a}_i', \hat{b}_i', b_i^*, \Gamma_3, \Gamma_4)$
2. **begin**
3. $p_v \leftarrow P_{\Gamma_3}^{(\tau_0)}(\hat{a}_i', v)$ for $v \in K_3$
4. $\hat{p}_v \leftarrow P_{\Gamma_4}^{(\tau_0)}(\hat{b}_i', v)$ for $v \in K_4$ (the distribution of b_i^*.)
5. $p_{\min} \leftarrow \min\{p_v : v \in K_3\}$
6. $\hat{p}_{\max} \leftarrow \max\{\hat{p}_v : v \in K_4\}$
7. $\hat{\ell} \leftarrow 0$
8.   **forever do**
9. $\hat{\ell} \leftarrow \hat{\ell} + 1$
10.   Pick a walk $\bar{W}_{\hat{\ell}}$ of length τ_0 in K_3 according to the distribution
               on trajectories, conditioned on start point $= \hat{a}_i'$
11.   Let $\hat{x}_{\hat{\ell}}$ be the terminal vertex of $\bar{W}_{\hat{\ell}}$
12.   With probability $\hat{p}_{\hat{x}_{\hat{\ell}}} p_{\min}/(p_{\hat{x}_{\hat{\ell}}} \hat{p}_{\max})$ accept $\bar{W}_{\hat{\ell}}$ and **exitloop**
13.   **od**
14.   **output** $\bar{W}_1, \bar{W}_2, \ldots, \bar{W}_{\hat{\ell}}$
15. **end** WALK1

Lemma 7 *In* WALK1, $\hat{x}_{\hat{\ell}}$ *is chosen according to the distribution* $\hat{p}$.

Proof: The probability s that a walk is accepted at the last step in the loop is given by

$$s = \sum_{v \in K_3} p_v \frac{\hat{p}_v p_{\min}}{p_v \hat{p}_{\max}} = \frac{p_{\min}}{\hat{p}_{\max}}. \tag{28}$$

18

(Observe that $\hat{p}_{\max} \geq 1/|K_4| = 1/|K_3| \geq p_{\min}$.) Thus

$$\mathbf{Pr}(\hat{x}_{\hat{\ell}} = v) = \sum_{k=0}^{\infty} (1-s)^k p_v \frac{\hat{p}_v p_{\min}}{p_v \hat{p}_{\max}} = \hat{p}_v. \tag{29}$$

□

Hence $\bar{W}_{\bar{r}}$ is a random walk to a vertex chosen with distribution $\hat{p}$. Furthermore, by (8),

$$s \geq \sigma = \frac{1}{9} \tag{30}$$

and therefore the expected number of walks generated is constant.

There is a minor problem in that we want to choose the endpoints before we do the walks. This leads to the algorithm WALK described before. We now turn to its analysis.

Lemma 8 *Suppose that b_i^* is chosen from K_3 with distribution $\hat{p}$. Then the walks $\bar{W}_1, \ldots, \bar{W}_{\hat{\ell}}$ in WALK1$(\hat{a}_i', \hat{b}_i', b_i^*, \Gamma_3, \Gamma_4)$, and the walks $\hat{W}_1, \ldots, \hat{W}_\ell$ in WALK$(\hat{a}_i', \hat{b}_i', b_i^*, \Gamma_3, \Gamma_4)$ have the same distribution.*

Proof: Note first from the proof of Lemma 7 that $\hat{\ell}$ and ℓ have the same truncated geometric distribution. Also, we have from Lemma 7 that $\hat{x}_{\hat{\ell}}$ and $x_\ell = b_i^*$ have the same distribution. Consider next that for $v_1, v_2, \ldots, v_i \in K_3$,

$$\mathbf{Pr}(\hat{x}_1 = v_1, \ldots, \hat{x}_i = v_i \text{ and } \hat{\ell} > i)$$

$$= \prod_{j=1}^{i} \left(\left(1 - \frac{\hat{p}_{v_j} p_{\min}}{p_{v_j} \hat{p}_{\max}} \right) p_{v_j} \right) = \prod_{j=1}^{i} \left(p_{v_j} - \frac{\hat{p}_{v_j} p_{\min}}{\hat{p}_{\max}} \right)$$

$$= (1-s)^i \prod_{j=1}^{i} \left(\frac{p_{v_j} - \hat{p}_{v_j} p_{\min}/\hat{p}_{\max}}{1-s} \right)$$

$$= \mathbf{Pr}(x_1 = v_1, \ldots, x_i = v_i \text{ and } \ell > i).$$

Thus $\hat{x}_1, \hat{x}_2, \ldots, \hat{x}_{\hat{\ell}}$ and $x_1, x_2, \ldots, x_\ell$ have the same distribution. Finally, the lemma follows from the fact that the distribution of $\bar{W}_j$ conditional on $\hat{x}_j = v$ is clearly equal to that of $\hat{W}_j$ conditional on $x_j = v$. □

Now we need to show that there are not many vertices in G_3 having degrees less that $r/10$ after the ℓ walks are deleted. We shall refer to these walks as the *i*th *bundle*, B_i. We shall closely follow the line of proof used in the analysis in Section 8.1, with all the events now referring to Γ_3 rather than Γ_2. As before, the proof reduces to showing that given $\mathcal{E}^{(I)}$, after the removal of the bundles $B_1, \ldots, B_i$ from G_3, the size of $\hat{V}_3$ (the set of vertices which are incident with at least $r/30$ edges of $B_1, \ldots, B_i$) is small.

Again, let S be a fixed vertex set of size $100 \ln n$ and q_k be the probability that S is visited by bundles $B_1, \ldots, B_i$ exactly k times. We now evaluate the number of visits to S. Assume first that $\hat{a}_i \in S$. In this case this we will distinguish between *free* visits and *start* visits. S undergoes ℓ visits, as the start point of all the walks in the bundle. All other visits to S are free visits. For $k \geq 2$, the $k - \ell$ free visits to S can be distributed among the ℓ walks in at most $\binom{k-1}{\ell-1}$ ways. We have

$$\mathbf{Pr}(k \text{ visits and } \hat{a}_i \in S) \leq \frac{100s}{n} \sum_{\ell=1}^{\infty} \binom{k-1}{\ell-1} \rho_S^{k-\ell} \sigma(1-\sigma)^{\ell-1} \leq \frac{C_4 s}{n} (1 - \sigma + \rho_S)^{k-1}$$

The term $\frac{100}{n}$ is arrived at as in (22).

Now consider the case $\hat{a}_1 \notin S$. Here we distinguish between *first* visits and free visits. A first visit is the first visit to S in a walk. Now we argue that there will be at least $k - \ell$ free visits and the probability of actually visiting S at least once in ℓ walks it at most $\frac{1000 s \ell \log_r n}{n}$. We obtain

$$\mathbf{Pr}(k \text{ visits and } \hat{a}_i \notin S) \leq \sum_{\ell=1}^{\infty} \frac{1000 s \ell \log_r n}{n} \binom{k-1}{\ell-1} \rho_S^{k-\ell} \sigma(1-\sigma)^{\ell-1} \leq \frac{C_5 s \log_r n}{n} (1-\sigma+\rho_S)^{k-1}.$$

So putting $\xi = (\sigma - \rho_S)/2 > 1/20$ we see that

$$\mathbf{E}((1+\xi)^{N_{i,S}} \leq 1 + \sum_{k \geq 1} \frac{C_6 s \log_r n}{n} (1-\xi)^{k-1}$$
$$\leq 1 + \frac{C_6 s \log_r n}{n}.$$

Arguing as from (25) to (26) we then obtain that

$$\mathbf{Pr}\left(\sum_{j=1}^{\kappa} Z_{j,S} \geq st \mid \mathcal{E}^{(I)}, \bar{\mathcal{P}} \right) \leq 2(1+\xi)^{s(-t/2+2\alpha r C_6)}.$$

We then take $t = r/30$ as before and obtain

$$\mathbf{Pr}(S \subseteq \hat{V}_3 \mid \mathcal{E}^{(i)}, \bar{\mathcal{P}}) \leq 2(1+\xi)^{-rs/120} < 2\epsilon_2.$$

We then argue, as before, that **whp** $|\hat{V}_3| \leq 3\epsilon_2 n$ and then use Lemma 3 to show that the $\lceil r/10 \rceil$-core of G_3, G_4 is of size at least $n - \epsilon_K n$.

8.3   Analysis of CONNECTBACK

We use the subroutine CONNECTBACK to connect vertices which are not in K back to K. In this subroutine we construct random walks using only uncovered edges. This effectively removes the possibility of long walks which go back and forward over the same edges.

Recall the configuration model of Section 5.1. Assuming (10) and

$$|V(\Gamma_5)| \geq n - ne^{-r/800} - o(n) \tag{31}$$

at each step of one of the random walks in CONNECTBACK, the probability that the next vertex of the walk is not in K is at most $3\epsilon_K$. The degrees of Γ_5 satisfy (7) and these walks model CONSTRUCT. (31) is justified by Lemma 4 and the fact that **whp** the $\lceil r/10 \rceil$-core of Γ_5 contains all but $o(n)$ vertices of G_5 – see the argument preceding (33) below. The main algorithm GENPATHS will call the subroutine CONNECTBACK at most 4κ times, so that the total number of random walks executed by CONNECTBACK is at most 4κ. Let W_i be the ith walk executed by CONNECTBACK. As before, fix v and let $N_{i,v}$ be the number of visits to $v \in G_5$ by the walk W_i, $Z_{i,v}$ be the number of edges incident with v which are traversed by W_i and let ℓ_i be the length of W_i. It is clear from the previous paragraph that $\mathbf{Pr}(\ell_i \geq (\ln n)^2) = O(n^{-A})$ for any constant $A > 0$. We can therefore consider a modified algorithm which terminates in failure if ever we find $\ell_i \geq (\ln n)^2$.

20

Note now that the probability that W_i starts at v is at most $10/n$ since either v is chosen uniformly from $K^{(0)}$ (Step 6 in GenPaths) or it is the endpoint of a long random walk (Steps 15, 17). Then

$$\mathbf{Pr}(N_{i,v} = k) = \sum_{\ell=k}^{(\ln n)^2} \mathbf{Pr}(N_{i,v} = k, \ell_i = \ell)$$

$$\leq \sum_{\ell=k}^{(\ln n)^2} (3\epsilon_K)^{\ell-1}(1 - 3\epsilon_K) \binom{\ell}{k} \left(\frac{10}{n}\right)^k \qquad (32)$$

$$\leq \frac{(3\epsilon_K)^{k-1}}{(1 - 3\epsilon_K)^k} \left(\frac{10}{n}\right)^k$$

$$\leq \frac{10}{n^k}$$

In (32) we have $(3\epsilon_K)^{\ell-1}(1 - 3\epsilon_K)$ as a bound for the probability that $\ell_i = \ell$. Then there are $\binom{\ell}{k}$ positions on the walk for the visits to v. Then $10/n$ bounds the probability that the walk makes the move to v at this point.

Note that if $v \in K$ then $N_{i,v} \leq 1$. Also, the upper bound we have calculated is valid, conditional on the previous history of the algorithm.

Suppose that there are $\kappa_1 \leq 4\kappa$ walks generated by ConnectBack overall. Then, for any $t > 0$

$$\mathbf{Pr}\left(\sum_{j=1}^{\kappa_1} N_{j,v} = t\right) \leq \sum_{\ell_1 + \cdots + \ell_{\kappa_1} = t} \prod_{j=1}^{\kappa_1} \left(\frac{10}{n}\right)^{\ell_j}$$

$$\leq \binom{\kappa + t - 1}{t} \left(\frac{10}{n}\right)^t$$

$$\leq \binom{4\kappa + t - 1}{t} \left(\frac{10}{n}\right)^t$$

$$\leq \left(\frac{40e\alpha r}{\log_r n}\right)^t$$

Let $\hat{V}_5$ be the set of vertices of G_5 which are incident with at least $r/30$ edges of $W_1, \ldots, W_i$. Then

$$\mathbf{Pr}(v \in \hat{V}_5) \leq \mathbf{Pr}\left(\sum_{j=1}^{\kappa} N_{j,v} \geq r/30\right)$$

$$\leq \sum_{t=r/30}^{2r/7} \mathbf{Pr}\left(\sum_{j=1}^{\kappa} N_{j,v} = t\right) \leq 2\left(\frac{40e\alpha r}{\log_r n}\right)^{r/30}$$

So if $\rho = n\left(\dfrac{40e\alpha r}{\log_r n}\right)^{r/40}$ then, by the Markov inequality,

$$\mathbf{Pr}(|\hat{V}_5| \geq \rho) \leq \left(\frac{40e\alpha r}{\log_r n}\right)^{r/120} \to 0.$$

If $\hat{V}_5$ is this small then Lemma 3 implies that the $\lceil r/10 \rceil$-core Γ_5 contains all but 2ρ vertices of G_5 (see the argument following (27)).

Thus
$$|L| \leq 2\rho \ \mathbf{whp} \tag{33}$$
as $L \subseteq G_5 \setminus \Gamma_5$.

9 Analysis of Phase 4

We join the pairs in L using the algorithm of [10]. The algorithm is capable of joining $\Omega(n/(\ln n)^c)$ distinct pairs for some constant $c > 0$, provided the graph has sufficent edge-expansion. Notice that $\tilde{a}_i, \tilde{b}_i$ are chosen as distinct vertices. For large r, G_6 will $\mathbf{whp}$ have ample expansion and 2ρ will be small enough for the algorithm.

References

[1] N. Alon, J. Spencer and P. Erdős, *The probabilistic method*, John Wiley and Sons, 1992.

[2] E. A. Bender and E. R. Canfield, *The asymptotic number of labelled graphs with given degree sequences*, Journal of Combinatorial Theory (A) 24, (1978) 296-307.

[3] B. Bollobás, *Random graphs*, Academic press, 1985.

[4] B. Bollobás, *A probabilistic proof of an asymptotic formula for the number of labelled regular graphs*, European Journal on Combinatorics 1, (1980) 311-316.

[5] B. Bollobás, *Martingale, isoperimetric inequalities and random graphs*, Coll. Math. Soc. J. Bolyai **52**, 113-139 (1987).

[6] B. Bollobás, T. I. Fenner, and A. M. Frieze, *Hamilton cycles in random graphs with minimal degree at least k*, in "A tribute to Paul Erdős," edited by A. Baker, B. Bollobás, and A. Hajnal (1990) 59-96.

[7] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*, North-Holland 1976.

[8] A. Z. Broder, A. M. Frieze, S. Suen, and E. Upfal, *Optimal Construction of Edge-Disjoint Paths in Random Graphs*, to appear.

[9] A. Z. Broder, A. M. Frieze, S. Suen, and E. Upfal, *Optimal construction of vertex-disjoint paths in random graphs*, in preparation.

[10] A. Z. Broder, A. M. Frieze, and E. Upfal, *Existence and construction of edge disjoint paths on expander graphs*, SIAM Journal on Computing 23 (1994) 976-989.

[11] A. Z. Broder, A. M. Frieze, and E. Upfal, *Static and dynamic path selection on expander graphs: a random walk approach*, to appear.

[12] P. Erdős and L. Lovász, Problems and results on 3-chromatic hypergraphs and some related questions. In A. Hajnal et al., editors, *Infinite and Finite Sets*, volume 11 of *Colloq. Math. Soc. J. Bolyai*, pages 609–627. North Holland, 1975.

[13] A.M.Frieze, Disjoint Paths in Expander Graphs via Random Walks: a Short Survey, to appear in Proceedings of Random '98.

[14] D. Gale, *A theorem on flows in networks*, Pacific Journal of Mathematics 7 (1957) 1073-1082.

[15] D. S. Hochbaum. *An exact sublinear algorithm for the max flow, vertex-disjoint paths, and communication problems on random graphs*, Operations Research 40 (1992) 923–935.

[16] J. Kahn and E. Szemerédi, *On the second eigenvalue in random regular graphs - Section 2*, Proceedings of the 21'st Annual ACM Symposium on Theory of Computing (1989) 587-598.

[17] J. Kleinberg and E. Tardos, *Approximations for the disjoint paths problem in high diameter planar networks*, Proceedings of the 27'th Annual ACM Symposium on Theory of Computing, (1995) 26-35.

[18] T.Leighton and S.Rao, Circuit switching: a multicommodity flow based approach, Proceedings of a Workshop on Randomized Parallel Computing 1996.

[19] T.Leighton, S.Rao and A.Srinivasan, Multi-commodity flow and circuit switching, *Proceedings of the Hawaii International Conference on System Sciences*, 1998.

[20] T.Leighton, S.Rao and A.Srinivasan, New algorithmic aspects of the local lemma with applications to partitioning and routing.

[21] C. J. H. McDiarmid, *On the method of bounded differences*, in Surveys in Combinatorics, Proceedings of the Twelfth British Combinatorial Conference (J. Siemons, Ed.), 1989, pp. 148-188.

[22] D. Peleg and E. Upfal, *Constructing disjoint paths on expander graphs*, Combinatorica 9, (1989) 289-313.

[23] N. Robertson and P. D. Seymour, *Graph minors-XIII: The disjoint paths problem*, to appear.

[24] E. Shamir and E. Upfal. *A fast construction of disjoint paths in networks*. Annals of Discrete Mathematics 24 (1985) 141–154.

[25] A. Sinclair and M. Jerrum, *Approximate counting, uniform generation, and rapidly mixing Markov chains*, Information and Computation 82 (1989) 93-133.

[26] L. G. Valiant and G. J. Brebner, *Universal schemes for parallel computation*, Proceedings of the 13'th Annual ACM Symposium on Theory of Computing (1981) 263-277.