Arc-Disjoint Paths in Expander Digraphs

Tom Bohman* and Alan Frieze†

Department of Mathematical Sciences,
Carnegie Mellon University
Pittsburgh PA 15213.

Abstract

Given a digraph $D = (V, A)$ and a set of κ pairs of vertices in V, we are interested in finding for each pair (x_i, y_i), a directed path connecting x_i to y_i, such that the set of κ paths so found is arc-disjoint. For arbitrary graphs the problem is $\mathcal{NP}$-complete, even for $\kappa = 2$.

We present a polynomial time randomized algorithm for finding arc-disjoint paths in an r-regular expander digraph D. We show that if D has sufficiently strong expansion properties and r is sufficiently large then *all* sets of $\kappa = \Omega(n/\log n)$ pairs of vertices can be joined. This is within a constant factor of best possible.

1 Introduction

Given a graph/digraph $D = (V, A)$ with n vertices and a set of κ pairs of vertices in V, we are interested in finding, for each pair (x_i, y_i), a directed path connecting x_i to y_i, such that the set of κ paths so found is edge/arc-disjoint. This is a classical problem in graph theory. See Frank [7] for a survey and Chapter 9.2 of the recent book on digraphs by Bang-Jensen and Gutin [2].

For undirected graphs, the related decision problem is in $\mathcal{P}$ for fixed κ – Robertson and Seymour [22], but is $\mathcal{NP}$-complete if κ is part of the input. For digraphs the situation is seemingly much worse. Fortune, Hopcroft and Wylie [6] showed that the related decision problem is $\mathcal{NP}$-complete, even when $\kappa = 2$.

For undirected graphs there have been positive results in the case of expanders. Peleg and Upfal [21] presented a polynomial time algorithm for the case where D 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. This result has been improved and extended by Broder, Frieze, and Upfal [4, 5], Frieze [8, 9], Leighton and Rao [17] and Leighton, Rao and Srinivasan [18, 19]. In particular Frieze [9] showed that if D has sufficiently strong *edge* expansion properties and r is sufficiently large then *all* sets of $\kappa = \Omega(n/\log n)$ pairs of vertices can be joined. This is within a constant factor of a simple upper bound. The purpose of this paper is to extend this result to digraphs.

In this paper we discuss r-regular digraphs. A digraph D is r-regular if every vertex has in-degree and out-degree r. Let d_μ be the median distance between pairs of vertices in the digraph D which has m arcs. Clearly, there exists a collection of $O(m/d_\mu)$ pairs of vertices that cannot be connected by arc-disjoint paths because such a collection of paths would require more arcs than all the arcs available. In the case of an r-regular expander, this absolute upper bound on κ is $O(n/\log n)$ (assuming r is independent of n). We show that if D has sufficiently strong *arc* expansion properties and r is sufficiently large then *all* sets of $\kappa = \Omega(n/\log n)$ pairs of vertices can be joined. This therefore, is within a constant factor of the optimum. The precise definition of "sufficiently strong" is given after the theorem.

Theorem 1 *Let $D = (V, A)$ be an n-vertex, r-regular digraph. Suppose that D is a sufficiently strong arc expander. Then there exist $\epsilon_1, \epsilon_2 > 0$ such that D has the following property: For all sets of pairs of vertices $\{(x_i, y_i) \mid i = 1, ..., \kappa\}$ satisfying:*

(i) $\kappa = \lceil \epsilon_1 rn/\log n \rceil$.
(ii) For each vertex v, $|\{i : x_i = v\}|, |\{i : y_i = v\}| \leq \epsilon_2 r$.

There exist arc-disjoint paths in D, each of length $O(\log n)$, joining x_i to y_i, for each $i = 1, 2, \ldots, \kappa$. Furthermore, there is a polynomial time randomized algorithm for constructing these paths. The constants ϵ_1, ϵ_2 depend only on certain expansion parameters α, β, γ defined below. They do not depend on n or r. (For example conditions (3) with $\epsilon = \alpha$, (4), (10) and (13) suffice).

Remark 1 *The algorithm is similar to the algorithm of [9]. The difficulty in moving from graphs to digraphs has been with*

*Supported in part by a grant from the NSA.
†Supported in part by NSF grant CCR-9818411.

that part of the algorithm for graphs which was based on the rapid mixing of a random walk on expanders. Random walks on digraphs are not necessarily time reversible and the steady state can be hard to determine. We therefore abandoned this approach and replaced it with a different random choice of path (in part, this random choice for digraphs uses the multi-commodity flow results of Leighton and Rao [16]).

It will be observed that this new algorithm can substitute for that given in [9].

Remark 2 *If D has sufficiently strong* vertex *expansion properties then we can take $\kappa = \lceil \epsilon_1 rn/\log_r n \rceil$ – see Remark 3 below.*

1.1 Preliminaries

In this section we state the definitions for the expanders we work with here, make some preliminary observations about such expanders, and make precise the notion of 'sufficiently strong' expansion needed for Theorem 1. We begin with some notational conventions.

Let $D = (V, A)$ be a digraph and let $n = |V|$. For $S, T \subseteq V$ let $A_D^+(S,T)$ be the set of arcs with tail in S and head in T; that is,

$$A^+(S,T) = A_D^+(S,T) = \{(u,v) \in A \mid u \in S, v \in T\},$$
$$\text{and } d^+(S,T) = |A^+(S,T)|.$$

We define $A^-(S,T)$ simillarly (i.e $A^-(S,T) = A^+(T,S)$). We set

$$A^+(S) = \{(u,v) \in A \mid u \in S, v \notin S\} \text{ and } d^+(S) = |A^+(S)|.$$

So, for example, we have $A^+(S) = A^+(S, V \setminus S)$. We define $A^-(S)$ and $d^-(S)$ analogously. Throughout the paper, when $\star$ is used as a subscript or superscript, it stands for + or -. We appreviate $A^*(\{v\}, T)$ to $A^\star(v,T)$. Thus, for $v \in V$, $d_D^+(v)$ and $d_D^-(v)$ denote the out-degree and in-degree of v in D.

We now have the notation necessary to introduce expansion. We define expanders in terms of arc expansion (a weaker property than vertex expansion). For $S \subseteq V$ let $\Phi_S^\star = d^\star(S)/|S|$. The (arc-)expansion $\Phi = \Phi(D)$ of D is defined by

$$\Phi = \min_{\substack{S \subseteq V \\ |S| \le n/2}} \min\{\Phi_S^+, \Phi_S^-\}.$$

A digraph $D = (V, A)$ is a θ-expander, if for every set $S \subseteq V$, $|S| \le n/2$, we have $d^\star(S) \ge \theta|S|$; in other words, D is a θ-expander if $\Phi(D) \ge \theta$. An r-regular digraph $D = (V, A)$ is called an (α, β, γ)-expander if for every set $S \subseteq V$

$$d^\star(S) \ge \begin{cases} (1-\alpha)r|S| & \text{if } |S| \le \gamma n \\ \beta r|S| & \text{if } \gamma n < |S| \le n/2 \end{cases}$$

We naturally assume that $\beta < 1 - \alpha$. By "sufficiently strong" in Theorem 1, we mean that β, γ are arbitrary and α is sufficiently small; in particular, we assume that conditions (3) with $\epsilon = \alpha$, (4), (10) and (13) hold. We also assume throughout that r and n are sufficiently large (but r is *not* a function of n). We have made no real attempt to optimize constants.

We conclude this section with some preliminary observations about expanders. Since $|A^\star(S,S)| + d^\star(S) = r|S|$ we see that, putting $in(S) = |A^+(S,S)| = |A^-(S,S)|$, in an (α, β, γ)-expander

$$in(S) \le \alpha r|S| \qquad \text{when } |S| \le \gamma n. \tag{1}$$

In particular random regular digraphs are usually (α, β, γ)-expanders. (See discussion in [4] for the corresponding notion in undirected graphs.) For a digraph $\Delta = (V', A')$ and a set $S \subseteq V'$ we define its out-neighbor set $N_\Delta^+(S)$, as

$$N_\Delta^+(S) = \{w \notin S : \exists v \in S \text{ such that } (v,w) \in A'\}.$$

Similarly, the in-neighbor set of S, $N_\Delta^-(S)$ is given by

$$N_\Delta^-(S) = \{w \notin S : \exists v \in S \text{ such that } (w,v) \in A'\}.$$

Lemma 1 *Suppose that $D = (V, A)$ is an (α, β, γ)-expander and that $D' = (V', A')$ is a sub-digraph of D of expansion at least θr where $\theta > \alpha$. Suppose $S \subseteq V'$. If $|S| \le \frac{\gamma\alpha}{\theta} n$ then*

$$|N_{D'}^\star(S)| \ge \frac{\theta - \alpha}{\alpha}|S|.$$

Proof: Suppose that $|S| \le \frac{\gamma\alpha}{\theta} n$ and $T = N_{D'}^\star(S)$ satisfies $|T| < \frac{\theta-\alpha}{\alpha}|S|$. Then

$$|S \cup T| \le \left(1 + \frac{\theta - \alpha}{\alpha}\right)|S| = \frac{\theta}{\alpha}|S| \le \gamma n.$$

But, $S \cup T$ contains at least

$$\theta r|S| > \theta r \left(1 + \frac{\theta - \alpha}{\alpha}\right)^{-1} |S \cup T| = \alpha r|S \cup T|$$

arcs, which contradicts (1). □

2 The algorithm

The input to our algorithm is a sufficiently strong (α, β, γ)-expander digraph D and a set of pairs of vertices $\{(x_i, y_i) \mid i = 1, ..., \kappa\}$ satisfying the premises of Theorem 1. The output is a set of κ arc-disjoint paths, $P_1, \ldots, P_\kappa$ such that P_i connects x_i to y_i.

The algorithm has three phases. We begin in Phase 0 by splitting our graph into 13 arc-disjoint expanders $D_i =$

(V, A_i), $1 \leq i \leq 13$. These graphs will be used for various purposes in Phases 1 and 2, and each path we construct will be a union of paths in the D_i's. Phase 1 consists mainly of applications of GenPaths, an algorithm that uses the expander property to naively connect pairs of vertices with paths of length $O(\log n)$ one at a time, deleting the arcs from a path as soon as it is used. Of course, this deletion of arcs may quickly destroy the expander property. To compensate for this problem, GenPaths 'shrinks' the expanders in which it finds paths. The real work of GenPaths is in keeping as many vertices as possible 'connected' to these shrinking expanders. Those pairs of vertices that are not connected with paths in Phase 1 (i.e. those pairs that contain a vertex whose connection with one of the 'shrinking expanders' is lost) are handled in Phase 2. There are $O\left(\frac{n}{\log^4 n}\right)$ such pairs. Loosely speaking, Phase 2 uses the multicommodity flow algorithm of Leighton and Rao to give a distribution on paths connecting the remaining pairs such that **whp** paths chosen at random with respect to this distribution are arc disjoint.

2.1 Phase 0.

We need an algorithm for splitting a (α, β, γ)-expander digraph into 13 expander digraphs. Algorithms for splitting undirected graphs are given in [4] or [10]. They are easily adapted to digraphs and we omit the details for lack of space. The expansion requirements for this algorithm are

$$\frac{r}{\log r} \geq 91\epsilon^{-2} \text{ and } \Phi \geq 65\epsilon^{-2} \log 2er, \tag{2}$$

Since the arc-expansion of a (α, β, γ)-expander is βr, we need

$$\beta \geq 65\epsilon^{-2} r^{-1} \log 2er. \tag{3}$$

In the full paper we prove:

Theorem 2 *Suppose that (2) holds and that G is an r-regular (α, β, γ)-expander, r constant. Then there is a randomised polynomial time algorithm which with probability at least $1 - \delta$ constructs $A_1, A_2, \ldots, A_{13}$ such that the arc-expansion Φ_i of $D_i = (V, A_i)$ satisfies*

$$\Phi_i \geq (1 - \epsilon)\frac{\Phi}{13} - (\alpha + 2\epsilon)\, r,$$

for $i = 1, 2, \ldots, 13$.

This theorem is only useful if Φ is at least a constant multiple of r and α is sufficiently small. This is the case discussed in this paper. The algorithm runs in $O(n^2 \ln n \log \delta^{-1})$ expected time.

We apply the algorithm of Theorem 2 with $\alpha = \epsilon$ and assume that

$$\beta > 156\alpha. \tag{4}$$

Setting

$$\beta_0 = \frac{\beta}{13} - 4\alpha > 8\alpha > 0, \tag{5}$$

each D_i satisfies

$$\Phi_i = \Phi(D_i) \geq \beta_0 r, , \qquad \text{and} \tag{6}$$

$$\beta_0 r \leq d_i^\star(v) < r, \qquad \text{for all } v \in V. \tag{7}$$

2.2 Phase 1.

Phase 1 uses expanders D_1 through D_6. The centerpiece of Phase 1 is the algorithm GenPaths that connects a large collection of *random* pairs of vertices in an expander with arc-disjoint paths. Since the pairs x_i, y_i are arbitrary, GenPaths cannot be applied to them directly; we must reduce the problem of connecting the x_i's to the y_i's to the problem of connecting random pairs.

In order to produce random pairs we introduce three random sets of κ vertices: $\tilde{X}$, $\tilde{Y}$ and Z. In the Initialization step of Phase 1, a network-flow technique is used to find a collection of arc-disjoint paths $\mathcal{P}^1 = \{P_i^1 : i = 1, \ldots, \kappa\}$ from $X = \{x_1, \ldots, x_\kappa\}$ to $\tilde{X}$ in expander D_1 such that the path P_i^1 starts at x_i for $i = 1, \ldots, \kappa$. It is important to note that we have no control over which element of $\tilde{X}$ is at the end of the path P_i^1 (but there will be one path ending at x for each $x \in \tilde{X}$). This network-flow technique is also used in the Initialization step of Phase 1 to find a collection of arc-disjoint paths $\mathcal{P}^6 = \{P_i^6 : i = 1, \ldots, \kappa\}$ from $\tilde{Y}$ to $Y = \{Y_1, \ldots, Y_\kappa\}$ in D_6 such that the endpoint of the path P_i^6 is y_i. After the Initialization step, we take a random ordering of $\tilde{X}$: $\tilde{X} = \{\tilde{x}_1, \ldots \tilde{x}_\kappa\}$. Furthermore, we order $\tilde{Y}$ so as to respect the pairing of $\tilde{X}$ and $\tilde{Y}$ that is inherited from the collections $\mathcal{P}^1$ and $\mathcal{P}^6$; that is, we set $\tilde{Y} = \{\tilde{y}_1, \ldots, \tilde{y}_\kappa\}$ so that if the endpoint of P_i^1 is $\tilde{x}_j$ then $\tilde{y}_j$ is the starting point of P_i^6. Thus, it remains to find a collection of arc-disjoint paths connecting the pairs $\{(\tilde{x}_i, \tilde{y}_i) : i = 1, \ldots, \kappa\}$. Unfortunately this sequence of pairs of vertices is *not* truly random. This is a consequence of the fact that the pairing between $\tilde{X}$ and $\tilde{Y}$ is determined by a deterministic process (i.e. knowledge of some of the *pairs* from this collection may bias the distribution on the unknown pairs). We overcome this problem by introducing the third random set $Z = \{z_1, \ldots, z_\kappa\}$. The sequences $\{(\tilde{x}_i, z_i) : i = 1, \ldots, \kappa\}$ and $\{(z_i, \tilde{y}_i) : i = 1, \ldots, \kappa\}$ are perfectly random sequences of pairs of vertices (so long as we view them separately). It remains to connect these two sequences of pairs of vertices with arc-disjoint paths. This is the work of the algorithm GenPaths.

The input to GenPaths is a pair of expanders, D_a and D_b, and a collection of pairs of vertices $\{(v_i, u_i) : i = 1, \ldots, \kappa\}$ that is generated uniformly at random. The output of GenPaths is a collection of arc-disjoint paths from v_i to u_i for

$i = 1, \ldots, \kappa$ that use only the arcs from D_a and D_b. We apply GENPATHS twice. For the first application we set $D_a = D_2$, $D_b = D_3$ and $\{(v_i, u_i) : i = 1, \ldots, \kappa\} = \{(\tilde{x}_i, z_i) : i = 1, \ldots, \kappa\}$. In the second application we set $D_a = D_4$, $D_b = D_5$ and $\{(v_i, u_i) : i = 1, \ldots, \kappa\} = \{(z_i, \tilde{y}_i) : i = 1, \ldots, \kappa\}$. Now, the expander D_a is used to connect the v_i's to the u_i's with *short* paths one at a time. In order to be sure that such paths exist we must be working with an expander. Therefore we delete some vertices in the course of the algorithm; in other words, this expander shrinks as the algorithm progresses. D_b is used to keep as many vertices as possible connected to the 'shrinking expander' contained in D_a. We should note that these connections also require that D_b be an expander. So, D_b also 'shrinks' in the course of the algorithm. The subroutines REMOVE and CONNECTBACK are used by GENPATHS.

2.2.1 Initialization

Let $\tilde{X}, \tilde{Y}$ be two randomly chosen κ-subsets of V. We begin by replacing the problem of finding paths from x_i to y_i by that of finding paths from a_i to b_i, where $a_i \in \tilde{X}$ and $b_i \in \tilde{Y}$. Let X denote the set $\{x_1, x_2, \ldots, x_\kappa\}$ and $Y = \{y_1, y_2, \ldots, y_\kappa\}$. We connect X to $\tilde{X}$ via arc-disjoint paths in the digraph D_1 using a network flow. We construct our network as follows

- Each directed arc of D_1 gets capacity 1.
- Each $v \in V$ becomes a source of capacity $|\{i : x_i = v\}|$ and each member of $\tilde{X}$ becomes a sink of capacity 1.

Then we find a flow from X to $\tilde{X}$ that satisfies all demands. Since the maximum flow has integer values, it decomposes naturally into $|X|$ arc-disjoint paths (together perhaps with some cycles). We connect Y to $\tilde{Y}$ by arc-disjoint paths in a similar manner using D_6

We now have a collection of arc-disjoint paths $\mathcal{P}^1 = \{P_i^1 : i = 1, \ldots \kappa\}$ such that the starting point of P_i^1 is x_i for $i = 1, \ldots, \kappa$ and each member of $\tilde{X}$ is the endpoint of exactly one path from $\mathcal{P}^1$. Furthermore, we have a collection of arc-disjoint paths $\mathcal{P}^6 = \{P_i^6 : i = 1, \ldots \kappa\}$ such that each member of $\tilde{Y}$ is the starting point of exactly one path from $\mathcal{P}^6$ and the end-point of path P_i^6 is y_i for $i = 1, \ldots, \kappa$.

2.2.2 Algorithm GENPATHS.

The aim of GENPATHS is to join v_i and u_i for $i = 1, 2, \ldots, \kappa$ by a *short* (i.e. of length $O(\log n)$) path in D_a. After constructing a path, we remove its arcs. It is important to ensure that short paths exist. Of course, this would not be a problem if we could ensure that D_a remains an expander throughout. We have to be satisfied with identifying a *dynamically changing large* subgraph $\Delta_a = (V_a, F_a)$ of D_a which is an expander. Initially $\Delta_a = D_a$, and V_a loses vertices as GENPATHS progresses. We ensure that Δ_a remains an expander by keeping the degrees of vertices in the Δ_a close to their degree in D_a. This may involve deleting some (low degree) vertices after the construction of a path. We use the routine REMOVE to do this.

If the proposed start vertex v of a walk on Δ_a does not lie in V_a then we try to *connect it back* to V_a by a path in D_b. The terminal endpoint of this walk is denoted by v'. We use a subroutine CONNECTBACK for this purpose. Similarly, the proposed end vertex u might not lie in V_a. In this case we use CONNECTBACK to find a path from some $u' \in V_a$ to u in D_b. We do not expect to succeed all the time and our failures are kept in a set L for treatment in Phase 2. The arcs in the paths generated by CONNECTBACK are deleted from D_b. Since CONNECTBACK requires that D_b is an expander we will also be working with a second 'shrinking expander' $\Delta_b = (V_b, F_b)$ contained in D_b. This shrinking expander will also be maintained by use of the subroutine REMOVE.

In the end, the path from v_i to u_i will be a concatenation of up to three separate paths. There will always be a path Q_i from D_a, and there may also be a short walk (or walks) from D_b provided by CONNECTBACK (these are denoted $W_i^{CB\rightarrow}$ and $W_i^{CB\leftarrow}$, respectively).

1. **Algorithm** GENPATHS
2. **begin**
3. $\Delta_i \leftarrow D_i$, $i = a, b$.
4. **for** $i = 1$ **to** κ **do**
5. **Execute** REMOVE(Δ_a)
6. **Execute** CONNECTBACK($V_a, v_i, \rightarrow, v_i', i, W_i^{CB\rightarrow}$)
7. **Execute** CONNECTBACK($V_a, u_i, \leftarrow, u_i', i, W_i^{CB\leftarrow}$)
8. **if** $i \notin L$ **then**
9. Construct a shortest path Q_i from v_i' to u_i' in Δ_a.
10. $P_i \leftarrow (W_i^{CB\rightarrow}, Q_i, W_i^{CB\leftarrow})$
11. $\Delta_a \leftarrow \Delta_a \setminus E(P_i)$
12. **fi**
13. **od**
14. **end** GENPATHS

2.2.3 Subroutine REMOVE

The purpose of REMOVE is to delete vertices which might prevent a digraph from being an expander. In the course of GENPATHS we apply REMOVE to Δ_a and Δ_b. In words, this simple algorithm iteratively removes those vertices whose in/out degree is less than the original in/out degree minus $\beta_0 r/2$. To be precise, at the end of the algorithm we have a graph $\Delta_t = (V_t, F_t)$ where $t \in \{a, b\}$ such that

$$v \in V_t \text{ implies } d^\star_{\Delta_t}(v) \geq d^\star_{D_t}(v) - \beta_0 r/2 \geq \beta_0 r/2. \qquad (8)$$

The final inequality in (8) follows from (7). It follows immediately from (8) that for $S \subseteq V_t$ we have

$$d^\star_{\Delta_t}(S) \geq d^\star_{D_t}(S) - \beta_0 r|S|/2 \geq (\Phi_t - \beta_0 r/2)|S|.$$

This implies that, provided neither V_a nor V_b become empty (an issue which we take up in the next section), both Δ_a and Δ_b are expanders throughout Phase 1:

$$\Phi_{\Delta_t} \geq \Phi_t - \beta_0 r/2 \geq \beta_0 r/2 \qquad \text{for } t = a, b. \tag{9}$$

1. **Algorithm** REMOVE(Δ_t)
2. **begin**
3. $B \leftarrow \{v \in V_t : d^+_{\Delta_t}(v) < d^+_{D_t}(v) - \beta_0 r/2$ or $d^-_{\Delta_t}(v) < d^-_{D_t}(v) - \beta_0 r/2\}$.
4. **if** $B \neq \emptyset$ **then**
5. $A \leftarrow V_t \setminus B$
6. $d \leftarrow \max_v\{ \max\{d^+_{D_t}(v) - d^+_{\Delta_t}(v, A), d^-_{D_t}(v) - d^-_{\Delta_t}(v, A)\} : v \in A\}$.
7. **while** $d > \beta_0 r/2$ **do**
8. $C \leftarrow \{w \in A : \max\{d^+_{D_t}(w) - d^+_{\Delta_t}(w, A), d^-_{D_t}(w) - d^-_{\Delta_t}(w, A)\} \geq \beta_0 r/2\}$
9. $B \leftarrow B \cup C$
10. $A \leftarrow A \setminus C$
11. $d \leftarrow \max_v\{ \max\{d^+_{D_t}(v) - d^+_{\Delta_t}(v, A), d^-_{D_t}(v) - d^-_{\Delta_t}(v, A)\} : v \in A\}$.
12. **od**
13. $V_t \leftarrow A$
14. **fi**
15. **end** REMOVE

2.2.4 Subroutine CONNECTBACK.

The purpose of CONNECTBACK is to connect a vertex z to V_a by means of a short walk in D_b. The direction of this walk is determined by the input dir. If dir $=\rightarrow$ then a path from z to V_a is required, and if dir $=\leftarrow$ then a path from V_a to z is needed. If $z \in V_a$ already then CONNECTBACK does nothing but relabel z as z'. Since $|V \setminus V_a|$ can be of order n (this is discussed below), we must maintain the expander property of Δ_b in order to find short connecting paths. Thus we apply REMOVE to Δ_b in the course of CONNECTBACK. Now, those pairs that contain a vertex that lies in $V \setminus (V_a \cup V_b)$ are passed to Phase 2 in the set L. Thus, the long term success of CONNECTBACK hinges on keeping $V \setminus (V_a \cup V_b)$ small. In fact, this can be viewed as the key point in all of GENPATHS.

We keep V_b large by ensuring that the paths we use in Δ_b are spread out; in other words, we avoid using too many paths through any one vertex. This is achieved **whp**; recall that the pairs of vertices $\{(v_i, u_i) : i = 1, \ldots, \kappa\}$ that are the input to GENPATHS are assumed to be generated uniformly at random. When a path is needed (i.e when $z \notin V_a$) CONNECTBACK constructs a collection $\mathcal{W}^{\text{dir}}$ of walks in Δ_b. This collection has the following properties.

1. If dir $=\rightarrow$ then every walk in $\mathcal{W}^{\text{dir}} = \mathcal{W}^{\rightarrow}$ is a walk from a distinct vertex in $V_b \setminus V_a$ to V_a.
2. If dir $=\leftarrow$ then every walk in $\mathcal{W}^{\text{dir}} = \mathcal{W}^{\leftarrow}$ is a walk from V_a to a distinct vertex in $V_b \setminus V_a$.
3. No path in $\mathcal{W}^{\text{dir}}$ is longer than $22 \log\log n$.
4. No vertex of D_b lies on more than $240(\log\log n)^2$ paths.

The set of start vertices of the walks in $\mathcal{W}^{\rightarrow}$ is denoted $S^{\rightarrow}_{CB}$ and the set of terminal vertices of the walks in $\mathcal{W}^{\leftarrow}$ is denoted $S^{\leftarrow}_{CB}$. Clearly, $S^{\text{dir}}_{CB} \subseteq V_b \setminus V_a$. The collection of walks will have the additional property

5. $|V \setminus (V_a \cup S^{\text{dir}}_{CB})| \leq \frac{n}{(\ln n)^4}$.

In words, condition 3 says that the paths in $\mathcal{W}^{\text{dir}}$ are short, condition 4 says that the paths are 'spread out' and condition 5 says that very few vertices are left out of the collection.

We emphasis that $\mathcal{W}^{\text{dir}}$ is constructed without use of any information about z. Therefore, z (which was a random vertex to begin with) can be viewed as a vertex chosen uniformly at random *after* the collection $\mathcal{W}^{\text{dir}}$ is constructed. Heuristically, we can think of $\mathcal{W}^{\rightarrow}$ and $\mathcal{W}^{\leftarrow}$ as collections of connecting paths that are updated whenever Δ_a or Δ_b 'shrink,' but we only 'look' at these collections when we need them.

If $z \in S^{\text{dir}}_{CB}$ then we connect z back to V_a by way of the unique path in $\mathcal{W}^{\text{dir}}$that begins at z (if dir $=\rightarrow$) or ends at z (if dir $=\leftarrow$). If z does not lie in S^{dir}_{CB} we put i into L, where $z = \tilde{x}_i$ or $\tilde{y}_i$. Arc disjoint paths for the pairs $(\tilde{x}_i, \tilde{y}_i), i \in L$ are found in Phase 2.

A network flow technique for the construction of the collection of paths $\mathcal{W}^{\text{dir}}$ follows from the proof of Lemma 3, which is given in Section 2.3.3 below.

1. **subroutine** CONNECTBACK($V_a, z, \text{dir}, z', i, W_{CB}$)
2. **begin**
3. **if** $z \in V_a$
4. **then** $z' \leftarrow z$
5. **else**
6. **Execute** REMOVE(Δ_b)
7. **Contruct** $\mathcal{W}^{\text{dir}}$ (see Lemma 3 for algorithm).
8. **if** $z \notin S^{\text{dir}}_{CB}$
9. **then** $L \leftarrow L \cup \{i\}$
10. **else**
11. $W_{CB} \leftarrow$ the unique path in $\mathcal{W}^{\text{dir}}$ with start/terminal vertex z
12. $z' \leftarrow$ terminal/start vertex of W_{CB}
13. $\Delta_b \leftarrow \Delta_b \setminus W_{CB}$
14. **fi**
15. **fi**
16. **end** CONNECTBACK

2.3 Analysis of Phase 1

There are three facts concerning Phase 1 that remain to be shown: that the flow needed in the Initialization exists, that

V_3 stays large and that at the end of Phase 1 **whp** we have $L = O\left(\frac{n}{(\log n)^4}\right)$.

2.3.1 Initialization

In this subsection we show that if (6) holds and r is sufficiently large then we can find arc-disjoint paths from $\{x_1, \ldots, x_\kappa\}$ to $\tilde{X}$ in D_1 and arc-disjoint paths from $\tilde{Y}$ to $\{y_1, \ldots, y_\kappa\}$ in D_6, for *any* choice of $x_1, \ldots, y_\kappa$ consistent with the premises of Theorem 1, and every choice for $\tilde{X}, \tilde{Y}$. We assume that we have

$$8\alpha > \epsilon_2 > r^{-1} \tag{10}$$

For $S \subseteq V$, let

$$\alpha(S) = \sum_{v \in S} |\{i : x_i = v\}| \qquad \text{and} \qquad \xi(S) = |S \cap \tilde{X}|.$$

It follows from a theorem of Gale [12] (see Bondy and Murty [3] Theorem 11.8) that if

$$d^+_{D_1}(S) \geq \xi(\bar{S}) - \alpha(\bar{S}), \qquad \forall S \subseteq V. \tag{11}$$

then there exists a flow in the network defined on D_1 such that exactly one unit of flow travels through each vertex in $\tilde{X}$ and the amount of flow traveling through each vertex $v \in \{x_1, \ldots, x_\kappa\}$ is $|\{i : x_i = v\}|$. In other words, (11) implies a successful run of Phase 2.

Now, if $|S| \leq n/2$ then, applying (5), (6) and (10), we have

$$d^+_{D_1}(S) \geq |S|\Phi_1 \geq |S|\beta_0 r \geq 8\alpha r|S| \geq \epsilon_2 r|S| \geq \alpha(S) - \xi(S) = \xi(\bar{S}) - \alpha(\bar{S}).$$

On the other hand, if $|S| > n/2$ then we have

$$d^+_{D_1}(S) = d^-_{D_1}(\bar{S}) \geq |\bar{S}|\Phi_1 \geq |\bar{S}|\beta_0 r \geq \epsilon_2 r|\bar{S}| \geq \xi(\bar{S}) - \alpha(\bar{S}).$$

Therefore, Phase 1 succeeds with respect to $X, \tilde{X}$. The same argument applies to $Y, \tilde{Y}$. To ensure these paths are of length $O(\log n)$ we can solve a minimum cost maximum flow problem as indicated in Kleinberg and Rubinfeld [13].

2.3.2 On the size of V_a

Lemma 2 *Throughout* GenPaths *we have*

$$|V_a| \geq (1 - \gamma_0)n$$

where

$$\gamma_0 = \frac{\beta_0 \gamma}{10}.$$

Proof: It follows from (9) that Δ_a is a $(\beta_0 r/2)$-expander throughout the execution of Phase 1. It follows from Lemma 1 that the diameter of Δ_a is always at most

$$\tau = \lceil 2\log_{1+\beta_0/2} n + \tfrac{4\gamma\alpha}{\log \beta_0 r} \rceil. \tag{12}$$

Indeed, consider breadth first search from some $v \in V_a$. Let $L_t, t \geq 0$ be the vertices at distance t from v. Lemma 1 implies that the cardinalities of the L_t grow at a rate at least $1+\beta_0/2$ until they reach size $(2\gamma\alpha/\beta_0)n$. The same will be true for breadth first search to a target vertex w. This accounts for the first term in (12). Once we have L_t reaches $(2\gamma\alpha/\beta_0)n$ we use that fact that going to the next level involves finding $\beta_0 r|L_t|/2$ "new arcs", at least until the size $n/2$ is reached. This accounts for the second term in (12).

Thus the total number of arcs in the paths that are removed from D_a is at most $\kappa\tau$. Let B be the set of vertices that are removed from Δ_a in the course of Phase 1, and let B_1 be the set of vertices in B incident with at least $\beta_0 r/4$ of the paths that are generated in D_a. We have

$$|B_1| \leq \frac{8\kappa\tau}{\beta_0 r} \leq \frac{\gamma_0 n}{2}$$

provided

$$\epsilon_1 \leq \frac{\beta_0^2 \gamma}{400} \log\left(1 + \frac{\beta_0}{2}\right). \tag{13}$$

where ϵ_1 is as in the statement of Theorem 1. Let $B_2 = B \setminus B_1$ (i.e. those vertices removed from Δ_a that lie on less than $\beta_0 r/4$ of the paths generated in D_a).

Assume for the sake of contradiction that $|B_2| > |B_1|$. Let B_3 be the *first* $|B_1|$ vertices of B_2 to join B. Note that the vertices in B_3 have a large degree to $B_1 \cup B_3$ (otherwise these vertices would remain in V_a). Applying (5) we have

$$in_{D_a}(B_1 \cup B_3) \geq \frac{\beta_0 r}{4}|B_3| = \frac{\beta_0 r}{8}|B_1 \cup B_3| > \alpha r|B_1 \cup B_3|.$$

This contradicts (1).

Therefore, $|B| = |B_1| + |B_2| \leq 2|B_1| \leq \gamma_0 n$. □

2.3.3 Analysis of ConnectBack

Of course, the first order of business here is to show how the collection of paths $\mathcal{W}^{\rm dir}$ is generated.

Lemma 3 *Suppose that $D = (V, A)$ is an (α, β, γ)-expander and that $D' = (V', A')$ is a sub-digraph of D of expansion at least θr where $\theta > 6\alpha$. Suppose that $S \subseteq V'$ and that $|S| \geq (1 - \frac{\gamma\alpha}{\theta})n$ and let $T = V' \setminus S$. Then there exists $T^\star \subseteq T$, such that D' contains a collection of walks $\mathcal{W}^{\rm dir} = \{W_v : v \in T^\star\}$ such that*

1. *for* dir $=\rightarrow$ / $\leftarrow$, v *is the start/terminal vertex of* W_v *for all* $v \in T^\star$,
2. *the terminal/start vertex of each* W_v *is in* S,
3. *each* W_v *is of length at most* $22 \log\log n$,
4. *no vertex of* D' *lies on more than* $240(\log\log n)^2$ *paths, and*
5. $|T \setminus T^\star| \leq \frac{n}{(\log n)^4}$.

Proof: Assume without loss of generality that dir $=\rightarrow$. For $i = 1, 2, \ldots$ let

$$T_i = \{v \in T : \ dist_{D'}(v, S) = i\},$$

and set $T_0 = S$. Since $N^+_{D'}(\cup_{k \geq i} T_k) \subseteq T_{i-1}$, it follows from Lemma 1 that we have

$$|T_{i-1}| \geq \zeta |T_i| \qquad \text{for } i \geq 1, \tag{14}$$

where $\zeta = \frac{\theta - \alpha}{\alpha} > 5$. Setting $i_0 = \lceil 11 \log\log n \rceil$ and $\hat{T} = \bigcup_{i \geq i_0} T_i$ It follows from (14) that we have

$$|\hat{T}| \leq \frac{n}{2(\log n)^4}. \tag{15}$$

Fix $1 \leq i < i_0$. We define a flow network $\mathcal{N}_i$. The vertex set of $\mathcal{N}_i$ is $\{s, t\} \cup \bigcup_{j=1}^{j_0} (C_j \cup C'_j)$ where C_1 and C'_1 are disjoint copies of T_i, for $2 \leq j \leq i$, C_j and C'_j are disjoint copies of $\bigcup_{\ell \geq i+1-j} T_\ell$ and, for $i < j \leq j_0 = 2i_0$, C_j and C'_j are disjoint copies of V'. The vertices s and t will be the source and sink, respectively, for the flow we introduce to $\mathcal{N}_i$. A vertex v in V' may appear many times in the vertex set of $\mathcal{N}_i$; a copy of v in C_j is denoted v_j, and a copy of v in C'_j is denoted v'_j. For ease of notation, we let ϕ be the map that takes the vertices of $\mathcal{N}_i$ to their corresponding vertices in V'. The arc-set of $\mathcal{N}_i$ is defined as follows. There is an arc from s to each vertex of C_1. Each $v \in S$ gives rise to arcs $(v'_j, t), i < j \leq j_0$. If $v'_j \in C'_j$ and $w_{j+1} \in C_{j+1}$ are such that (v, w) is an arc of D' then (v'_j, w_{j+1}) is an arc of $\mathcal{N}_i$. All arcs described so far have infinite capacity. In addition there are arcs (v_j, v'_j) of unit capacity defining a perfect matching between C_j and C'_j for $1 \leq j \leq j_0$.

Claim 1 *$\mathcal{N}_i$ contains an $s - t$ flow of value at least $|T_i| - \frac{n}{(\log n)^5}$.*

We first show the Lemma follows from Claim 1. The flow given by Claim 1 defines paths in D' from all but at most $\frac{n}{(\log n)^5}$ vertices of T_i to S, each of length at most j_0. No vertex of V' can be on more than j_0 paths since each visit to v uses a (v_j, v'_j) arc for some j. Repeating this construction for $i = 1, 2, \ldots, i_0$ we find paths for all but a set $\tilde{T}$ of at most $i_0 \frac{n}{(\log n)^5} \leq \frac{n}{2(\log n)^4}$ vertices and no vertex can be on more than $i_0 j_0$ paths. Putting $T^\star = T \setminus (\hat{T} \cup \tilde{T})$ and using (15) gives us the lemma.

It only remains to prove Claim 1. Of course, we do this via the max-flow min-cut theorem. Consider a cut $Z \dot{\cup} \bar{Z}$ of $\mathcal{N}_i$ where the vertex set Z contains s but not t. Let $A_j = Z \cap C_j$, $B_j = C_j \setminus A_j$, $A'_j = Z \cap C'_j$ and $B'_j = C'_j \setminus A'_j$ for $j = 1, 2, \ldots, j_0$. Assume for the sake of contradiction that the capacity of this cut is less than $|C_1| - \frac{n}{(\log n)^5}$. It follows from this assumption that the cut contains no infinite capacity arcs and therefore

$$A_1 = C_1, \qquad \phi(A_{j+1}) \supseteq N^+_{D'}(\phi(A'_j)) \text{ for all } j, \quad \text{and } \phi(A'_j) \cap S = \emptyset \text{ for all } j. \tag{16}$$

The capacity of the cut is

$$|B'_1| + \sum_{j=2}^{j_0} |\phi(A_j) \cap \phi(B'_j)|. \tag{17}$$

The third condition of (16) implies that for all j, $|A'_j| < \frac{\gamma\alpha}{\theta} n$. It then follows from Lemma 1 and the second condition of (16) that we have $|A_{j+1}| \geq \zeta |A'_j|$. This implies that for all j we have either

$$|A'_{j+1}| \geq \frac{1}{2}|A_{j+1}| \geq \frac{\zeta}{2}|A'_j| \tag{18}$$

or

$$|\phi(A_{j+1}) \cap \phi(B'_{j+1})| \geq \frac{1}{2}|A_{j+1}| \geq \frac{\zeta}{2}|A'_j|. \tag{19}$$

Now, if $|A'_1| \leq \frac{n}{(\log n)^5}$ then $|B'_1| \geq |C_1| - \frac{n}{(\log n)^5}$, which contradicts our initial assumption. On the other hand, if $|A'_1| > \frac{n}{(\log n)^5}$ then, since $(\frac{\zeta}{2})^{j_0 - 1} > (\log n)^5$, condition (18) cannot always hold. Let $j_1 \geq 1$ be the first j for which (19) holds. We have

$$|\phi(A_{j_1+1}) \cap \phi(B'_{j_1+1})| \geq (\zeta/2)^{j_1} |A'_1| \geq |A'_1|.$$

The capacity of the cut is at least $|B'_1| + |A'_1| = |C_1|$. This is a contradiction. □

To get the collection of paths needed for ConnectBack we apply Lemma 3 with $D' = \Delta_b, V' = V_b, S = V_a \cap V_b$. So, for example, we have $S^{\text{dir}}_{CB} = T^\star$. It remains to show that **whp** we have $|L| \leq \frac{n}{(\log n)^4}$.

Note that Lemma 3 can only be applied if $V_a \cap V_b$ is large. However, by applying the proof of Lemma 2 and the fact that the paths generated by ConnectBack are short (length at most $22 \log\log n$) we see that $|V_b| = n - O(n \log\log n / \log n))$ throughout and this is sufficient. However, we shall see that V_b is **whp** larger than this. This is where we use the fact

that the paths in $\mathcal{W}^{\rm dir}$ are spread out (i.e. the fact that there are at most $240(\log\log n)^2$ paths in $\mathcal{W}^{\rm dir}$ through any one vertex).

It follows from this fact that the probability that that an arbitrary vertex w is on the path $W_i^{CB\rm dir}$ at most $240(\log\log n)^2/n$. It follows that we have

$$\begin{aligned}\mathbf{Pr}\left(\left|\left\{i : w \in W_i^{CB\rightarrow}\right\}\right| + \left|\left\{i : w \in W_i^{CB\leftarrow}\right\}\right| \geq 20\right)\\ \leq \mathbf{Pr}(B(2\kappa, 240(\log\log n)^2/n) \geq 20)\\ \leq \binom{2\kappa}{20}\left(\frac{240(\log\log n)^2}{n}\right)^{20}\\ = o((\log n)^{-19}).\end{aligned} \tag{20}$$

Let B be the set of vertices which are removed from Δ_b by applications of REMOVE. Let X_1 be the set of vertices in B that are on at least 20 of the paths $W_i^{CB\rightarrow}, W_i^{CB\leftarrow}$. Note that X_1 contains the set B_1 introduced in the proof of Lemma 2 (B_1 is the collection of vertices taken out of V_b by REMOVE that are on many of the paths). It follows from (20) and Markov's inequality that **whp** we have $|X_1| \leq n/(\log n)^{18}$. Now, B (which contains $V \setminus V_b$ at every step) consists of B_1 together with extra vertices deleted by REMOVE. In total this will be at most $2|B_1|$ vertices removed by the argument of Lemma 2, following (13). Thus, $|B| \leq 2|B_1| \leq 2|X_1|$. Therefore

$$|B| \leq \frac{n}{(\log n)^{18}} \tag{21}$$

whp.

Now, a failure (i.e. the index i joining the set L) can occur in one of two ways. On one hand we have a failure if either v_i or u_i does not lie in $V_a \cup V_b$, and on the other hand a failure results when $u_i \in V_b \setminus (V_a \cup S_{CB}^{\rightarrow})$ or $v_i \in V_b \setminus (V_a \cup S_{CB}^{\leftarrow})$. It follows from (21) that the total number of failures of the first type is **whp** at most $\frac{n}{(\log n)^{18}}$. For failures of the second type we note that $v_1, \ldots, v_\kappa$ form a random sequence of size $o(n)$. The probability that a particular vertex gives a failure of the second kind is at most

$$\frac{|S_{CB}^{\rm dir}|}{n} = O\left(\frac{1}{(\log n)^4}\right).$$

Therefore, **whp** the total number of failures of the second kind is $O(n/(\log n)^4)$.

Remark 3 *Suppose G has the following vertex expansion property for small sets: $S \subseteq V$, $|S| \leq \frac{\gamma}{r}n$ implies that $|N_D^\star(S)| \geq (1-\alpha)r|S|$. It is possible to split D so that each subgraph D_i and small S satisfies $|N_{D_i}^\star(S)| \geq \frac{1-30\alpha}{13}r|S|$. Then the shortest paths in Δ_a will be of length $O(\log_r n)$ and the claim in Remark 2 will follow.*

2.4 Phase 2.

The set of pairs $\{(\tilde{x}_i, \tilde{y}_i) : i \in L\}$ have not yet been connected by paths. We have seen that **whp** the number of such pairs, $|L|$, is at most $O(n/(\log n)^4)$. These pairs are dealt with by the algorithm described below which uses digraphs D_7–D_{13}.

The heart of the algorithm is a randomized method (based on a multicommodity flow result of Leighton and Rao [16]) for connecting pairs of vertices with arc disjoint paths which works **whp** when the collection of pairs is generated uniformly at random. So, as in Phase 1, some preliminary steps must be taken in order to reduce the problem of connecting an arbitrary set of pairs of vertices with arc-disjoint paths to the problem of connecting a random collection of pairs with arc-disjoint paths. We proceed directly to a description of the algorithm. Let $m = |L|$ and $\lambda = \lceil \log n \rceil$.

We begin by 'amplifying' each start vertex $\tilde{x}_i, i \in L$ and each end vertex $\tilde{y}_i, i \in L$ to a collection of λ vertices. This process occurs in steps **(a)** and **(b)**.

(a) In this step, we choose a collection of vertices $w_j, 1 \leq j \leq 2m$, and a collection of sets of vertices $W_j, 1 \leq j \leq 2m$, such that for $1 \leq j \leq 2m$,

- **(i)** $w_j \in W_j$,
- **(ii)** $|W_j| = \lambda$,
- **(iii)** The sets W_j are pairwise disjoint, and
- **(iv)** D_9 contains an arborescence with vertex set W_j and root w_j. For $1 \leq j \leq m$ this arborescence is directed away from the root and for $m+1 \leq j \leq 2m$ this arborescence is directed toward the root.

Following [17], we find these arborescences by partitioning large arborescences of D_9. We begin with a rooted spanning arborescence T with the property that all arcs are directed away from the root. We generate $W_1, \ldots, W_m$ greedily from D_9, removing an arborescence from T once it is used. Of course, this process will divide T into a number of components. However, since the maximum degree of D_9 is r, the number of components produced in this process is at most $r\lambda m = O\left(\frac{nr}{(\log n)^3}\right)$. Since any tree having at least λ vertices contains a subtree having exactly λ vertices, we will always be able to find the needed arborescences.

We then apply REMOVE to D_9 less the vertex set $\cup_{i=1}^m W_i$ to produce an expander D_9'. It follows from the proof of Lemma 2 that D_9' has $n(1-o(1))$ vertices. We repeat the process described above (this time using D_9' and an arborescense directed towards the roots) to produce $W_{m+1}, \ldots, W_{2m}$.

(b) Let $S_X = \{\tilde{x}_i : i \in L\}$ and $S_Y = \{\tilde{y}_i : i \in L\}$ denote the sets of vertices that need to be joined. Use a network flow algorithm (analogous to what is given in the Initialization step of Phase 1) in D_7 to connect in an arbitrary manner the vertices of S_X to to $W_X = \{w_1, \ldots, w_m\}$ by m arc disjoint paths. Using the same network flow algorithm in D_8, connect in an arbitrary manner the vertices of $W_Y = \{w_{m+1}, \ldots, w_{2m}\}$ to S_Y by m arc disjoint paths. The expansion properties of D_7 and D_8 ensure that such paths always exist (as we saw in the Initialization step of Phase 1).

Let $\hat{x}_k$ (resp. $\hat{y}_k$) denote the vertex in W_X that was connected to the end-point $\tilde{x}_k$ (resp. $\tilde{y}_k$). Our problem is now to find arc-disjoint paths joining $\hat{x}_k$ to $\hat{y}_k$ for $1 \le k \le m$. If w_t has been renamed as $\hat{x}_k$ (resp. $\hat{y}_k$) then rename the elements of W_t as $\hat{x}_{k,\ell}$, (resp. $\hat{y}_{k,\ell}$) $1 \le \ell \le \lambda$.

(c) Choose $\xi_j, 1 \le j \le \lambda m$, and $\eta_j, 1 \le j \le \lambda m$, uniformly at random from V. Using a network flow algorithm (as in (b)) connect $\{\hat{x}_{k,\ell} : 1 \le k \le m, 1 \le \ell \le \lambda\}$ to $\{\xi_j : 1 \le j \le \lambda m\}$ by arc-disjoint paths in D_{10}. Similarly, connect $\{\eta_j : 1 \le j \le \lambda m\}$ to $\{\hat{y}_{k,\ell} : 1 \le k \le m, 1 \le \ell \le \lambda\}$ by arc-disjoint paths in D_{13}. Rename the other endpoint of the path starting at $\hat{x}_{k,\ell}$ (resp. ending at $\hat{y}_{k,\ell}$) as $x^*_{k,\ell}$ (resp. $y^*_{k,\ell}$).

(d) Choose $z^*_{k,\ell}, 1 \le k \le m, 1 \le \ell \le \lambda$ uniformly at random, with replacement. Now, it is important to note that the pairs $x^*_{k,\ell}, z^*_{k,\ell}$ and the pairs $z^*_{k,\ell}, y^*_{k,\ell}$ are (when viewed separately) perfectly random. We are using the same 'trick' that we used in Phase 1 for replacing the problem of connecting arbitrary pairs to the problem of connecting random pairs.

The paths between pairs of the form $x^*_{k,\ell}, z^*_{k,\ell}$ and between pairs of the form $z^*_{k,\ell}, y^*_{k,\ell}$ are generated at random. Using the multicommodity flow algorithm of Leighton and Rao [16], find a collection of paths $P_{u,v;\theta}, u \neq v \in V, 1 \le \theta \le \nu_{u,v}$ where each $P_{u,v;\theta}$ is a path in D_{12} from u to v. These are the flow paths given by Theorem 18 in [16] ($\nu_{u,v}$ is simply the number of flow paths we have for the pairs u, v). Let $\mathcal{P}_{u,v} = \{P_{u,v;\theta} : 1 \le \theta \le \nu_{u,v}\}$. For each $P_{u,v;\theta}$ we will have a flow value $f_{u,v;\theta} > 0$ and we let $F_{u,v} = \sum_{\theta=1}^{\nu_{u,v}} f_{u,v;\theta}$. The properties we require are that

(P1) For all arcs e of D_{10},

$$\sum_{(u,v,\theta):\, e \in P_{u,v;\theta}} f_{u,v;\theta} \le 1.$$

(P2)

$$F_{u,v} \ge \frac{c_2}{n \log n}$$

for some absolute constant $c_2 > 0$.

(P3) The length of each path $P_{u,v;\theta}$ is at most $\lambda_1 = c_1 \log_r n$ for some absolute constant $c_1 > 0$.

For $u, v \in V$ let $P_{u,v}$ be the probability distribution over $\mathcal{P}_{u,v}$ where $P_{u,v}(P_{u,v;\theta}) = f(u,v,\theta)/F(u,v)$. Then for each k, ℓ choose $W'_{k,\ell}$ randomly from $\mathcal{P}_{x^*_{k,\ell}, z^*_{k,\ell}}$ using the distribution $P_{x^*_{k,\ell}, z^*_{k,\ell}}$ to select the path.

Let B'_k denote the *bundle* of paths $\{W'_{k,\ell}, 1 \le \ell \le \lambda\}$. Carry out the same construction in D_{13} and construct a bundle of paths $B''_k = \{W''_{k,\ell}, 1 \le \ell \le \lambda\}$ where $W''_{k,\ell}$ is a path from $z^*_{k,\ell}$ to $y^*_{k,\ell}$.

Let $\pi_0 = \max_e \mathbf{Pr}(e \in P)$ where P is a path chosen by (i) randomly choosing endpoints u, v and then (ii) choosing $P \in \mathcal{P}_{u,v}$ according to the distribution $P_{u,v}$. We have

$$\pi_0 = \max_e \sum_{P_{u,v;\theta} \ni e} \frac{1}{n^2} \cdot \frac{f(u,v,\theta)}{F(u,v)} \le \frac{1}{n^2} \frac{n \log n}{c_2} = \frac{\log n}{c_2 n}.$$

We say that $W'_{k,\ell}$ is *bad* if there exists $k' \neq k$ such that $W'_{k,\ell}$ shares an arc with a walk in a bundle $B'_{k'}$.

Now, suppose the bundles in the set $\{B_j : j \neq k\}$ are *fixed*. The collection of paths involved in these bundles gives at most $m\lambda\lambda_1$ arcs. Thus, the probability that $W_{k,l}$ is bad, conditioning on what happens outside the bundle B_k, is at most

$$\pi_0 m \lambda \lambda_1 = O\left(\frac{1}{\log n}\right).$$

We say that index k is bad if either B'_k or B''_k contain more than $\lambda/3$ bad walks. If index k is not bad then we can find a walk from $x^*_{k,\ell}$ to $y^*_{k,\ell}$ through $x^*_{k,\ell}$ for some ℓ which is arc disjoint from all other walks. This gives a walk

$$x_k - \tilde{x}_k - \hat{x}_k - \hat{x}_{k,\ell} - x^*_{k,\ell} - z^*_{k,\ell} - y^*_{k,\ell} - \hat{y}_{k,\ell} - \hat{y}_k - \tilde{y}_k - y_k,$$

which is arc-disjoint from all other such walks.

The probability that index k is bad is at most

$$2\,\mathbf{Pr}(B(\lambda, O((\log n)^{-1}) \ge \lambda/3) = O(n^{-2}).$$

So with probability 1-o(1) there are no bad indices.

□

References

[1] N. Alon and J.H. Spencer, *The Probabilistic Method,* Wiley, 1992.

[2] J. Bang-jensen and G. Gutin, Digraphs: Theory, Algorithms and Applications, Springer-Verlag, London 2001.

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

[4] 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.

[5] A. Z. Broder, A. M. Frieze, and E. Upfal, *Existence and construction of edge low congestion paths on expander graphs*, Random Structures and Algorithms 14 (1999) 87-109.

[6] S. Fortune, J.E. Hopcroft and J. Wyllie, *The directed subgraph homeomorphism problem*, Theoretical Computer Science 10 (1980) 111-121.

[7] A. Frank, *Disjoint paths in rectilinear grids*, Combinatorica 2, (1982) 361-371.

[8] A.M. Frieze, *Disjoint Paths in Expander Graphs via Random Walks: a Short Survey*, Proceedings of Random '98, Lecture Notes in Computer Science 1518 (1998) Springer, 1-14.

[9] A.M. Frieze, *Edge disjoint paths in expander graphs*, *SIAM Journal on Computing* 30 (2001) 1790-1801.

[10] A.M.Frieze and M.Molloy, *Splitting an expander graph*, Journal of Algorithms 33 (1999) 166-172.

[11] A.M. Frieze and L. Zhao, *Edge disjoint paths in random regular graphs*, Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (1999) 291-299 .

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

[13] J. Kleinberg and R. Rubinfeld, *Short paths in expander graphs*, Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science, (1996) 86-95.

[14] 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.

[15] D.E. Knuth, The art of computer programming, Volume 1, Fundamental Algorithms, Addison-Wesley, 1968.

[16] T. Leighton and S. Rao, *Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms*, Journal of the Association for Computing Machinery 46 (1999) 787-832.

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

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

[19] T. Leighton, S. Rao and A. Srinivasan, *New algorithmic aspects of the local lemma with applications to partitioning and routing*, Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (1999) 643-652..

[20] A. Lubotsky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica 8 (1988) 261-277.

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

[22] N. Robertson and P. D. Seymour, *Graph minors-XIII: The disjoint paths problem*, *Journal of Combinatorial Theory Series B* 63 (1995) 65-110.

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

[24] D. Wagner and K. Weihe, *A linear time algorithm for edge-disjoint paths in planar graphs*, Proceedings of the First European Symposium on Algorithms (ESA '93) Lecture Notes in Computer Science 726, Springer-Verlag (1992) 384-395.

[25] X. Zhou, S. Tamura and T. Nishizeki, *Finding edge-disjoint paths in partial k-trees*, Algorithmica 26 (2000) 3-30.