Cycle partitions of regular graphs

Vytautas Gruslys* Shoham Letzter†

Abstract

Magnant and Martin conjectured that the vertex set of any d-regular graph G on n vertices can be partitioned into $n/(d+1)$ paths (there exists a simple construction showing that this bound would be best possible). We prove this conjecture when $d = \Omega(n)$, improving a result of Han, who showed that in this range almost all vertices of G can be covered by $n/(d+1)+1$ vertex-disjoint paths. In fact, our proof gives a partition of $V(G)$ into cycles. We also show that, if $d = \Omega(n)$ and G is bipartite, then $V(G)$ can be partitioned into $n/(2d)$ paths (this bound in tight for bipartite graphs).

1 Introduction

Dirac's classical result states that every graph on $n \geq 3$ vertices with minimum degree at least $n/2$ contains a Hamilton cycle. This minimum degree condition is best possible, as there is no Hamilton cycle in the almost balanced complete bipartite graph $K_{\lfloor (n-1)/2 \rfloor, \lceil (n+1)/2 \rceil}$ nor in the graph obtained by overlapping two cliques, $K_{\lfloor (n+1)/2 \rfloor}$ and $K_{\lceil (n+1)/2 \rceil}$, at a single vertex. While this means that Dirac's result cannot be extended to general graphs with minimum degree lower than $n/2$, such an extension may be possible if certain natural conditions are imposed on the graph. A very nice conjecture, posed independently by Bollobás [2] and Häggkvist (see [12]), stated that if $d \geq n/(t+1)$ then every t-connected d-regular graph on n vertices is Hamiltonian. It is indeed natural to require the graph be regular so that imbalanced complete bipartite graphs are ruled out. Note that the case $t = 1$ follows directly from Dirac's theorem.

The conjecture of Bollobás and Häggkvist has been resolved. The case $t = 2$ was proved by Jackson [12], following partial results of Nash-Williams [27], Erdős and Hobbs [7], and Bollobás and Hobbs [3]. Jackson's result was strengthened slightly by Hilbig [11], who showed that there

*Email: `vytautas.gruslys@gmail.com`.

†Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK. Email: `s.letzter@ucl.ac.uk`. Research was supported by Dr. Max Rössler, by the Walter Haefner Foundation and by the ETH Zurich Foundation.

are only two extremal examples (that is, 2-connected d-regular graphs with $d \geq n/3 - 1$ and no Hamilton cycle), namely, the Petersen graph and the graph obtained by replacing one vertex of the Petersen graph by a triangle. Following a number of partial results by Fan [8], Jung [15], Li and Zhu [28], Broersma, van den Heuvel, Jackson and Veldman [4], and Jackson, Li and Zhu [14], the case $t = 3$ was recently proved by Kühn, Lo, Osthus and Staden in two papers where they first obtained an asymptotic result [18] and then the exact result (for large n) [19]. This completed the picture regarding the Bollobás and Häggkvist conjecture, since the conjecture is false for $t \geq 4$, as was shown by Jung [15] and by Jackson, Li and Zhu [14].

A different direction was suggested by Enomoto, Kaneko and Tuza [6]: rather than finding one Hamilton cycle, they were interested in finding a small collection of cycles that covers the vertex set. More precisely, they conjectured that the vertices of any n-vertex graph with minimum degree at least d can be covered by at most $(n-1)/d$ cycles, where edges are considered to be cycles on two vertices. Note that the case where $d = n/2$ is exactly Dirac's theorem. The bound $(n-1)/d$ cannot be meaningfully lowered, since at least $\lfloor (n-1)/d \rfloor$ cycles are needed to cover the vertices of $K_{n-d,d}$ or of the graph obtained by taking one vertex of full degree and covering the other $n-1$ vertices by $\lfloor (n-1)/d \rfloor$ disjoint cliques, each of order at least d. Following progress by Enomoto, Kaneko and Tuza [6] and Kouider [16], this conjecture was proved by Kouider and Lonc [17] and, much later, but independently, by Balogh, Mousset and Skokan [1] for $d = \Omega(n)$.

What if the cycles in the conjecture of Enomoto, Kaneko and Tuza are required to be vertex-disjoint? In this case imbalanced bipartite graphs are again problematic, and so it makes sense to consider regular graphs. Magnant and Martin [26] conjectured that the vertices of any n-vertex d-regular graph can be covered by at most $n/(d+1)$ vertex-disjoint paths; this bound is tight as can be seen by taking a disjoint union of cliques of order $d+1$ (and, possibly, a larger d-regular graph on the remaining $d+1$ to $2d+1$ vertices). They proved this conjecture for $d \leq 5$ and Han [10] proved that, if $d = \Omega(n)$, then all but $o(n)$ vertices can be covered by at most $n/(d+1)+1$ paths. It does not seem critical that Magnant and Martin stated their conjecture for paths and not for cycles, because (at least in dense graphs) typical methods that give path partitions tend to give cycle partitions just as well. In this paper we prove Magnant and Martin's conjecture when $d = \Omega(n)$ and, indeed, our proof gives a partition into cycles.

Theorem 1. *For every $c_{\min} > 0$ there exists n_0 such that if $n \geq n_0$, $d \geq c_{\min} n$ and G is a d-regular graph on n vertices, then $V(G)$ can be partitioned into at most $n/(d+1)$ cycles.*

We also obtain an analogous result for bipartite graphs, but this time we only establish the existence of a path partition. The reason why our proof does not work for cycles seems to be technical rather than essential: we do believe that the same approach can give a proof for cycles, provided that some of our lemmas, including the main lemma of Section 5, are expanded with further technical conditions. However, to maintain the readability of this paper, we do not pursue this marginally stronger result.

Theorem 2. *For every $c_{\min} > 0$ there exists n_0 such that if $n \geq n_0$, $d \geq c_{\min} n$ and G is a d-regular bipartite graph on n vertices, then $V(G)$ can be partitioned into at most $n/(2d)$ paths.*

Theorem 2 improves a result of Han [10], who proved that all but $o(n)$ vertices can be covered by at most $n/(2d)$ vertex-disjoint paths. The bound $n/(2d)$ can be seen to be tight by taking a disjoint union of $\lfloor n/(2d) \rfloor$ $K_{d,d}$'s (possibly, replacing one of them by a slightly bigger d-regular bipartite graph, making sure that exactly n vertices are used).

In the following section we outline the proofs and the structure of the rest of the paper.

2 Overview

2.1 Outline of the proof

Our plan for proving Theorem 1 is as follows. (The proof of Theorem 2 is similar and, in fact, slightly simpler.) First, we partition the vertices into a small number of parts, which we call clusters, that are well-connected and such that there are few edges with ends in different clusters (this is made precise in Lemma 3). Kühn, Lo, Osthus and Staden [18, 19] used a similar partition. Moreover, the clusters in our partition can be shown to be *robust expanders*, a term that was introduced by Kühn, Osthus and Treglown [21] and has since proved to be very useful (see, for instance, [20, 22, 23]).

We zoom in on each cluster: ideally, we would like each one of them to be Hamiltonian and remain Hamiltonian after the removal of any small set of vertices. We establish this fact about all clusters that cannot be made bipartite by removing a small number of edges. However, the statement may fail for other clusters; for example, an imbalanced bipartite graph may appear as a cluster, and it is certainly not Hamiltonian. For clusters that are almost bipartite we establish a more technical statement: they become Hamiltonian after the removal of any small set of vertices that balances its two sides. This is done in Lemma 4, whose proof follows relatively easily from a result in [21].

Up to this point our argument mostly follows the strategy in [19]. Our main new ideas are in the proof of the next lemma, Lemma 5, in which we construct a small linear forest whose removal balances the clusters that are almost bipartite. A similar linear forest was constructed by Kühn, Lo, Osthus and Staden [18, 19]. However, their approach was more ad hoc and relied on the number of clusters being small (namely, at most five), whereas here this number can be arbitrarily large.

Upon the removal of the interior vertices of this linear forest, the clusters become Hamiltonian; in them we pick Hamilton paths that attach to the leaves of the linear forest. This ensures that the paths in the linear forest can be concatenated with the Hamilton paths in the clusters. The result is a small family of vertex-disjoint paths – containing no more paths than there are clusters – that covers the whole graph. By doing this step carefully, we ensure that each path in the family starts and ends at adjacent vertices, which means that this family is in fact a family of cycles.

2.2 Key lemmas

In this subsection we give some definitions and state Lemmas 3 to 5.

Here and later, we freely use standard definitions in graph theory: $e(H)$ denotes the number of edges of a graph H and, for disjoint sets $X, Y \subset V(H)$, we denote by $H[X, Y]$ the graph with vertex set $X \cup Y$ whose edges are the $X - Y$ edges of H (that is, those edges of H with one end in X and one in Y). Let G be a graph on n vertices. A *cut* of a set $A \subset V(G)$ is a partition $\{X, Y\}$ of A, where X and Y are both non-empty. We say that a cut $\{X, Y\}$ is *α-sparse* if $e(G[X, Y]) \leq \alpha|X||Y|$. We say that a set $A \subset V(G)$ is *α-almost-bipartite* if there exists a partition $\{X, Y\}$ of A such that $G[A]$ has at most αn^2 edges that are not $X - Y$ edges. Otherwise, we say that A is *α-far-from-bipartite.*

The following lemma, which is very similar to a result from [18], partitions the vertices of G into a small number of well-behaved sets, which we call *clusters.*

Lemma 3. *Let $c_{\min} \in (0, 1)$ and $n_0 \in \mathbb{N}$ be such that $1/n_0 \ll c_{\min}$. Let G be a d-regular graph on n vertices, where $n \geq n_0$ and $d \geq c_{\min} n$. Then there exist parameters $r \leq 1/c_{\min}$ and $\eta, \beta, \gamma, \zeta, \delta$, where $1/n_0 \ll \eta \ll \beta \ll \gamma \ll \zeta \ll \delta \ll c_{\min}$, and a partition $\{A_1, \ldots, A_r\}$ of $V(G)$ into non-empty sets satisfying the following properties:*

(a) *G has at most ηn^2 edges with ends in different A_i's;*

(b) *for each $i \in [r]$, the minimum degree of $G[A_i]$ is at least δn;*

(c) *for each $i \in [r]$, A_i has no ζ-sparse cuts;*

(d) *for each $i \in [r]$, A_i is either β-almost-bipartite or γ-far-from-bipartite.*

The meaning of the symbol $\ll$ requires some clarification. Every expression of the form $a \ll b$ should be read as 'a is much less than b'. Formally, it means that $a < \Phi(b)$ where $\Phi : (0, 1] \to (0, 1]$ is a hidden increasing function associated to that particular expression. The hidden functions depend only on the constant $c_{\min}$, and they can be worked out by carefully following the forthcoming arguments. We shall not mention these function again; instead, we shall implicitly assume that, as the variable approaches 0, they decrease sufficiently fast to make our calculations work.

We remark that the statement of Lemma 3 is somewhat unusual in that, given n_0, $c_{\min}$ and G as in the lemma, the conclusion holds for *some* choice of parameters $\eta, \beta, \gamma, \zeta, \delta$, with $1/n_0 \ll \eta \ll \beta \ll \gamma \ll \zeta \ll \delta \ll c_{\min}$, but not for every choice of such parameters. In particular, the correct choice for parameters depends on the graph G.

Given a graph G on n vertices, a set $A \subset V(G)$ is called *ξ-Hamiltonian* if, for any subset W of size at most ξn and any pair of distinct vertices $x, y \in A \setminus W$, there is a Hamilton path in $G[A \setminus W]$ with ends x, y. Given a partition $\{X, Y\}$ of A, we say that A is *ξ-weakly-Hamiltonian with respect to* $\{X, Y\}$ if, for any subset W of size at most ξn that satisfies $|X \setminus W| = |Y \setminus W|$ and any vertices $x \in X \setminus W$, $y \in Y \setminus W$, there is a Hamilton path in $G[A \setminus W]$ with ends x, y.

The following lemma shows that clusters are Hamiltonian if they are far from being bipartite and weakly-Hamiltonian if they are almost bipartite.

Lemma 4. *Let $c_{\min} \in (0, 1)$ and $n \in \mathbb{N}$, and let $\eta, \beta, \xi, \gamma, \zeta, \delta$ be real numbers satisfying $1/n \ll \eta \ll \beta \ll \xi \ll \gamma \ll \zeta \ll \delta \ll c_{\min}$. Let G be a d-regular graph on n vertices, where $d \geq c_{\min} n$, and suppose that $A \subset V(G)$ satisfies the following properties.*

(a) *there are at most ηn^2 edges in G with exactly one end in A;*

(b) *$G[A]$ has minimum degree at least δn;*

(c) *A has no ζ-sparse cuts;*

(d) *A is either β-almost-bipartite or γ-far-from-bipartite.*

If A is γ-far-from-bipartite, then A is ξ-Hamiltonian; if A is β-almost-bipartite, then it is ξ-weakly-Hamiltonian with respect to any partition $\{X, Y\}$ of A that maximises the number of X – Y edges.

When presented with a partition into well-behaved clusters, the next lemma produces a collection of vertex-disjoint paths that balances the clusters.

Lemma 5. *Let $c_{\min} \in (0, 1)$ and $n \in \mathbb{N}$, and let $\eta, \beta, \xi, \gamma, \zeta, \delta$ be real numbers satisfying $1/n \ll \eta \ll \beta \ll \xi \ll \gamma \ll \zeta \ll \delta \ll c_{\min}$. Let G be a d-regular graph on n vertices, where $d \geq c_{\min} n$, and let $\{A_1, \ldots, A_r\}$ be a partition of $V(G)$ with properties* (a) *to* (d) *in Lemma 3, where $r \leq \lceil 1/c_{\min} \rceil$. For each $i \in [r]$ such that A_i is β-almost-bipartite, let $\{X_i, Y_i\}$ be a partition of A_i that maximises the number of X_i – Y_i edges. Then there is a linear forest $H \subset G$ with the following properties:*

(a) *$|H| \leq \xi n$;*

(b) *H has no isolated vertices;*

(c) *for each $i \in [r]$, A_i contains either two or zero leaves of H;*

(d) *for each $i \in [r]$ such that A_i is β-almost-bipartite, either A_i contains no leaves of H, or X_i and Y_i each contain exactly one leaf of H;*

(e) *for each $i \in [r]$ such that A_i is β-almost-bipartite, $|X_i \setminus V(H)| = |Y_i \setminus V(H)|$.*

2.3 Proof of the main result

We now complete the proof of Theorem 1, using Lemmas 3 to 5. The proof mostly puts the three lemmas together, but we need to work a bit to get the exact right number of cycles. The lemmas themselves will be proved in forthcoming sections.

Proof of Theorem 1. Let $c_{\min} > 0$; we assume, without loss of generality, that $1/c_{\min} \in \mathbb{N}$. Let $n_0 \in \mathbb{N}$ satisfy $1/n_0 \ll c_{\min}$, and let G be a d-regular graph on n vertices, where $n \geq n_0$ and $d \geq c_{\min} n$. Let $\{A_1, \ldots, A_r\}$ be a partition of $V(G)$ produced by Lemma 3; this partition comes with parameters $1/n_0 \ll \eta \ll \beta \ll \gamma \ll \zeta \ll \delta \ll c_{\min}$. Set $l = \lfloor n/(d+1) \rfloor$ and let α be such that $\delta \ll \alpha \ll c_{\min}$.

For the moment, we fix a single index $i \in [r]$. By property (b) in Lemma 3, $|A_i| \geq \delta n$. Hence, by property (a), there exists a vertex $u \in A_i$ incident with at most $(\eta/\delta)n$ edges of G that leave A_i. Therefore, u has at least $d - (\eta/\delta)n$ neighbours in A_i, and so $|A_i| \geq d - (\eta/\delta)n \geq d\,(1 - \eta/(\delta c_{\min})) \geq (1-\alpha)d$ (using $d \geq c_{\min} n$ and $\eta \ll \beta \ll \alpha \ll c_{\min}$). More can be said if A_i is β-almost-bipartite. In such case we fix a partition $\{X_i, Y_i\}$ of A_i that maximises the number of X_i–Y_i edges in G. In particular, $G[X_i, Y_i]$ can be obtained from $G[A_i]$ by removing at most βn^2 edges. Similarly to the argument above, there exists a vertex in A_i, say in X_i, with at least $d - ((\eta + 2\beta)/\delta)\,n$ neighbours in $G[X_i, Y_i]$, which means that $|Y_i| \geq d(1 - (3\beta/\delta c_{\min})) \geq (1-\alpha)d$. Therefore, some vertex in Y_i has at least $d - (\eta + 2\beta)n^2/(1-\alpha)d$ neighbours in $G[X_i, Y_i]$, which implies that $|X_i| \geq d\left(1 - 3\beta/(1-\alpha)c_{\min}^2\right) \geq (1-\alpha)d$. We conclude that $|A_i| \geq (1-\alpha)d$ in general and $|A_i| \geq 2(1-\alpha)d$ if A_i is β-almost-bipartite.

Since $i \in [r]$ was arbitrary, we have $n \geq (r+s)(1-\alpha)d$, where s is the number of β-almost-bipartite A_i's. It follows that $r + s \leq l + 1$: this can be seen by bounding the difference

$$r + s - l \leq \left\lfloor \frac{n}{(1-\alpha)d} - \left\lfloor \frac{n}{d+1} \right\rfloor \right\rfloor \leq \left\lfloor \frac{\alpha d n + n}{(d+1)d(1-\alpha)} + 1 \right\rfloor \leq \left\lfloor \frac{\alpha + 1/n}{(c_{\min})^2(1-\alpha)} + 1 \right\rfloor = 1.$$

The rest of the proof splits into two cases: when $r \leq l$ and when $r = l + 1$. We first deal with the former case, which is critical, but easy to resolve using Lemma 5. We fix an arbitrary number ξ such that $\beta \ll \xi \ll \gamma$. Let H be a linear forest as produced by Lemma 5 (for each β-almost-bipartite A_i we use the partition $\{X_i, Y_i\}$ that was defined earlier in the proof), and we denote by I the set of internal vertices of H. For each $i \in [r]$, if A_i contains two leaves of H, then let x_i, y_i be those leaves. Otherwise, let $x_i, y_i \in A_i \setminus I$ be arbitrary adjacent vertices. Recall that $|I| \leq \xi n$ by property (a) in Lemma 5. We make two further observations if A_i is β-almost-bipartite. First, property (d) in Lemma 5 enables us to assume that $x_i \in X_i$ and $y_i \in Y_i$. Second, properties (d) and (e) in Lemma 5 imply that $|X_i \setminus I| = |Y_i \setminus I|$. Now, we apply Lemma 4 and conclude that, regardless of A_i being β-almost-bipartite or γ-far-from-bipartite, $G[A_i \setminus I]$ has a Hamilton path with ends x_i, y_i. We take these paths for all $i \in [r]$: some of them can be concatenated with the path components of H, while the others have adjacent ends and so can be completed into cycles. The result is a family of cycles that partitions $V(G)$. Note that the number of cycles in this family does not exceed the number of clusters, which is $r \leq l$.

We move on to the next case, that is, when $r = l + 1$. This immediately implies that $s = 0$, meaning that all A_i's are γ-far-from-bipartite. Suppose that there is a matching of size 2 between two distinct clusters A_i, A_j, and denote its edges by $x_i x_j$ and $y_i y_j$, where $x_i, y_i \in A_i$ and $x_j, y_j \in A_j$. By Lemma 4, for each $k \in \{i, j\}$ there is a Hamilton path in $G[A_k]$ with ends x_k and y_k. Together with the edges $x_i y_i$ and $x_j y_j$, we obtain a cycle whose vertex set is $A_i \cup A_j$. For every $k \neq i, j$, we use Lemma 4 again to find a cycle with vertex set A_k. In total we obtain a partition of the vertices into l cycles.

Now, let us assume for contradiction that there are no two distinct clusters with a matching of size 2 between them (i.e. the non-isolated vertices of $G[A_i, A_j]$ form a star for every $i \neq j$). We construct an auxiliary digraph H on vertices $V(G)$, whose arcs correspond to edges of G that join separate clusters. More precisely, for any distinct $i, j \in [r]$ such that G contains $A_i - A_j$ we do the following. If there is exactly one $A_i - A_j$ edge xy, we add both xy and yx to H. Otherwise, since $G[A_i, A_j]$ is a star with at least two edges, there is a unique vertex $a \in A_i \cup A_j$ such that all $A_i - A_j$ edges are incident with a. Add to H all $A_i - A_j$ edges as arcs directed towards a.

In order to complete the proof, we reach a contradiction by a double-counting argument. Intuitively, the structure of H suggests that there are few edges with ends in distinct parts A_i, but the assumption that there are $r = l + 1$ parts A_i implies that there are relatively many such edges. Fix $i \in [r]$. Let a_i denote the number of arcs in H that enter A_i and let b_i denote the number of arcs that leave A_i. Our most immediate aim is to establish the inequality

$$b_i \geq (d - 2l)(d + 1 - |A_i|) + a_i - 2l^2. \tag{1}$$

To this aim, we first observe that $|A_i| \geq d - l$, or else in G every vertex of A_i would send at least $l + 1$ edges to the other clusters. By the pigeonhole principle, at least two of these edges would end in the same cluster, and hence, again by the pigeonhole principle, there would be a cluster A_j, $j \neq i$, such that at least $|A_i|/l \geq 2$ vertices in A_i send at least two edges to A_j. However, this would contradict the assumption that there is no matching of size 2 between any two clusters. Furthermore, for any $j \neq i$, all arcs of H that go from A_j to A_i have the same head. Therefore, there are at least $d - 2l$ vertices in A_i of zero in-degree in H. We pick a set $Z \subset A_i$ consisting of exactly $d - 2l$ such vertices.

We write m for the number of $(A_i \setminus Z) - Z$ edges missing from G and denote the number of vertices in A_i of non-zero in-degree in H by k. We already know that $k \leq l$. In G, these vertices together send at least a_i edges outside of A_i, and so they send at most $kd - a_i$ edges to Z. Therefore, $m \geq k(d - 2l) - (kd - a_i) \geq a_i - 2l^2$. Since $\sum_{z \in Z} |N_G(z) \cap A_i| \leq |Z|(|A_i| - 1) - m$, there are at least

$|Z|(d+1-|A_i|)+m \geq (d-2l)(d+1-|A_i|)+a_i-2l^2$ edges from Z to $V(G)\setminus A_i$. They become arcs of H directed away from A_i, proving inequality (1).

Summing inequality (1) over $i \in [r]$, we get

$$0=\sum_{i=1}^{r}(b_i-a_i) \geq (d-2l)\left((d+1)r-n\right)-2l^2 r.$$

Since $r=\lfloor n/(d+1)\rfloor+1 > n/(d+1)$, we have $(d+1)r-n \geq 1$, and hence the right hand side of the inequality above is at least $c_{\min}n-2l-2l^2(l+1)>0$, giving a contradiction. □

In the proof of our main theorem, which we have just completed, we partitioned $V(G)$ into at most $l=\lfloor n/(d+1)\rfloor$ cycles. This proof can be tweaked so that exactly l cycles are guaranteed: if the original proof produces $l'<l$ cycles, then before invoking Lemma 4 to find Hamilton paths in the clusters, we can first take aside $l-l'$ very short cycles in one of the clusters (short cycles exist in clusters by Proposition 6).

If, instead of cycles, we wanted to partition $V(G)$ into (at most) l paths, then the analysis of the case $r=l+1$ in the proof of Theorem 1 would be simpler. Indeed, instead of finding a matching of size 2 between two clusters it would be enough to find a single edge.

2.4 Proof of the bipartite analogue

We now prove Theorem 2, which is the bipartite analogue of our main result. As long as we have Lemmas 3 to 5 at our disposal, the proof is straightforward, but, again, some care is needed to obtain the exactly tight bound.

Proof of Theorem 2. Let $c_{\min}$ be such that $1/c_{\min} \in \mathbb{N}$, and suppose that $1/n_0 \ll c_{\min}$. Let G be a bipartite d-regular graph on n vertices, where $n \geq n_0$ and $d \geq c_{\min}n$. Let X, Y be the vertex classes of G and write $l=\lfloor n/(2d)\rfloor$. Let $\{A_1,\dots,A_r\}$ be a partition of $V(G)$ as given by Lemma 3, where $1/n_0 \ll \eta \ll \zeta \ll \delta \ll c_{\min}$ are the corresponding parameters (β and γ do not play a role here as the graph is bipartite). The argument that applied to β-almost-bipartite clusters in the proof of Theorem 1 also works here and it shows that given α that satisfies $\delta \ll \alpha \ll c_{\min}$, we have $|A_i| \geq 2d(1-\alpha)$ for all $i \in [r]$. Therefore,

$$r-l \leq \left\lfloor \frac{n}{2d(1-\alpha)}-\frac{n}{2d}+1\right\rfloor \leq \left\lfloor \frac{\alpha}{2c_{\min}(1-\alpha)}+1\right\rfloor=1.$$

Let ξ be a parameter satisfying $\eta \ll \xi \ll \zeta$ and for each $i \in [r]$ fix the partition $\{X_i, Y_i\}$ for A_i, where $X_i=A_i\cap X$, $Y_i=A_i\cap Y$. Let H be a linear forest as given by Lemma 5. Precisely as in the proof of Theorem 1, by concatenating components of H and paths in the clusters, we partition

$V(G)$ into at most r cycles. Furthermore, if at least one component of H has ends in separate clusters, then the partition contains at most $r-1$ cycles. Therefore, we may assume that $r = l+1$ and both ends of each component of H are in the same cluster, as otherwise we are done.

Now, suppose that H has an edge uv with ends in separate clusters, say, $u \in X_1$, $v \in Y_2$. Let P be the component of H that contains uv, and let x_1, y_1 be the ends of P in X_1, Y_1, respectively (both parts of A_1 contain an end of P by property (d) in Lemma 5). We write P_u, P_v for the two paths comprising $P \setminus \{uv\}$, where P_u contains u and P_v contains v (P_u and/or P_v is a single vertex if u and/or v is an end of P). We select a vertex $x_2 \in X_2$ in the following way: if H has a component with ends in A_2, then we let x_2 be its end in X_2; otherwise, we pick x_2 arbitrarily. Note that $|(X_1 \setminus V(H)) \cup \{x_1\}| = |(Y_1 \setminus V(H)) \cup \{y_1\}|$ by property (e) in Lemma 5, and hence Lemma 4 produces a path with ends x_1, y_1 that spans $(A_1 \setminus V(H)) \cup \{x_1, y_1\}$. Similarly, there is a path spanning $(A_2 \setminus V(H)) \cup \{x_2, v\}$ that has ends x_2, v. Let P^* be the concatenation of P_u with the newly produced path between x_1, y_1, with P_v, with the newly produced path between v, x_2 and, if it exists, with the component of H whose one end is x_2. We observe that P^* is a path that covers $A_1 \cup A_2$ except for the vertices that appear in components of H with ends in clusters other than A_1, A_2. Outside of $A_1 \cup A_2$, P^* covers the vertices contained in components of H with ends in A_1, A_2. We deal with the clusters A_i for $i \geq 3$ in the same way as in the proof of Theorem 1. This gives a partition of $V(G)$ into the path P^* and at most $r-2 = l-1$ cycles, proving the result.

The final case to consider is when $r = l+1$ and H has no edges with ends in separate clusters. By property (d) in Lemma 5, every component of H covers the same number of vertices in both parts of the graph. Therefore, for each $i \in [r]$, $|X_i| = |X_i \setminus V(H)| + |X_i \cap V(H)| = |Y_i \setminus V(H)| + |Y_i \cap V(H)| = |Y_i|$. In other words, each cluster is balanced. Since $r > n/(2d)$, we may assume that $|A_1| < 2d$, and so $|X_1| = |Y_1| < d$. By the regularity of G, there exists an edge $uv \in E(G)$ with $u \in X_1$ and v not in Y_1. Say, $v \in Y_2$. By Lemma 4, for each $i \in [r]$ we may pick a path P_i spanning A_i, where u is an end of P_1 and v is an end of P_2. This gives a partition of $V(G)$ into $r-1 = l$ paths, namely, $P_1 uv P_2$, $P_3, \ldots, P_r$. $\square$

We remark that a possible strategy for proving a stronger version of Theorem 2 that establishes a partition of $V(G)$ into at most $\lfloor n/(2d) \rfloor$ cycles may revolve around moving a small number of vertices from some clusters to others, so that the clusters still satisfy the assumptions of Lemma 4, but the balancing linear forest now has a component with ends in separate clusters. We believe that we have a good idea on how such a proof would work – a more technical version of Lemma 5 is needed – but we decided not to pursue it.

2.5 Structure of the paper

We prove Lemmas 3 to 5 in Sections 3 to 5, respectively, and conclude the paper in Section 6 with closing remarks and open problems.

3 Partitioning the graph into well-behaved clusters

In this section we prove Lemma 3. This lemma is very similar to Theorem 3.1 in [18][1]. Nevertheless, as the proof in [18] is quite long, we give a proof here.

Proof of Lemma 3. Set $r_0 := \lceil 1/c_{\min} \rceil$ and fix positive constants n_0 and $\eta_1, \ldots, \eta_{r_0}$ that satisfy the hierarchy $1/n_0 \ll \eta_1 \ll \cdots \ll \eta_{r_0} \ll c_{\min}$. Let G be a cn-regular graph on $n \geq n_0$ vertices, where $c \geq c_{\min}$. We shall define a list $\mathcal{P}_1, \ldots, \mathcal{P}_r$, where $1 \leq r \leq r_0$, of increasingly refined partitions of $V(G)$ such that the following properties hold for each $i \in [r]$:

(i) $\mathcal{P}_i$ is a partition of $V(G)$ consisting of i non-empty parts;

(ii) if $i \geq 2$, then $\mathcal{P}_i$ is obtained by splitting one part of $\mathcal{P}_{i-1}$ into two;

(iii) G has at most $4\sqrt{\eta_{i-1}}n^2$ edges with ends in different parts of $\mathcal{P}_i$ (where $\eta_0 = 0$ by convention);

(iv) for every $A \in \mathcal{P}_i$, the minimum degree of $G[A]$ is at least $3^{-(i-1)}cn$;

(v) every part of $\mathcal{P}_r$ has no η_r-sparse cuts.

Let $\mathcal{P}_1 = \{V(G)\}$ and note that $\mathcal{P}_1$ trivially satisfies the first four conditions. Assuming that $\mathcal{P}_i$ is defined, we define $\mathcal{P}_{i+1}$ in the following way. If every part of $\mathcal{P}_i$ has no η_i-sparse cut, then we set $r = i$ and stop the process. Otherwise, we pick a part $A \in \mathcal{P}_i$ that has an η_i-sparse cut $\{A_1, A_2\}$. In A_1, we let A_1' be the set of vertices that have at most $\sqrt{\eta_i}n$ neighbours in A_2; similarly, we denote by A_2' the set of vertices in A_2 that have at most $\sqrt{\eta_i}n$ neighbours in A_1. Since $\{A_1, A_2\}$ is an η_i-cut of A, we have $\sqrt{\eta_i}n|A \setminus (A_1' \cup A_2')| \leq 2\eta_i n^2$, and hence $|A \setminus (A_1' \cup A_2')| \leq 2\sqrt{\eta_i}n$. Since every vertex in A has at least $3^{-(i-1)}cn$ neighbours in A and since all but at most $2\sqrt{\eta_i}n < 3^{-i}cn$ of them are in $A_1' \cup A_2'$, every vertex in A has at least $3^{-i}cn$ neighbours in A_j' for some $j \in \{1,2\}$. In particular, $G[A_1']$ and $G[A_2']$ both have minimum degree at least $3^{-i}cn$. Furthermore, we can partition $A \setminus (A_1' \cup A_2')$ into sets A_1'', A_2'' where for each $j \in \{1,2\}$ every vertex in A_j'' has at least $3^{-i}cn$ neighbours in A_j'. We define $\mathcal{P}_{i+1}$ by replacing the part A in $\mathcal{P}_i$ with two parts $A_1' \cup A_1''$ and $A_2' \cup A_2''$. It is clear that $\mathcal{P}_{i+1}$ satisfies properties (i), (ii) and (iv).

We now prove that $\mathcal{P}_{i+1}$ satisfies property (iii), provided that $i \leq r_0$ (we will show in the next paragraph that the process in fact terminates at some $\mathcal{P}_r$ with $r \leq r_0$). The number of edges between $A_1' \cup A_1''$ and $A_2' \cup A_2''$ is at most $\eta_i n^2 + |A_1'' \cup A_2''|cn \leq (\eta_i + 2\sqrt{\eta_i})n^2 \leq 3\sqrt{\eta_i}n^2$. Hence, by property (iii) of $\mathcal{P}_i$ and by the assumption that $\eta_{i-1} \ll \eta_i$, the number of edges between the parts of $\mathcal{P}_{i+1}$ is at most $(4\sqrt{\eta_{i-1}} + 3\sqrt{\eta_i})n^2 \leq 4\sqrt{\eta_i}n^2$, as desired.

If the process does not terminate for any $i \leq r_0$, then we create a partition $\mathcal{P}_{r_0+1}$ that satisfies properties (i) to (iv). We will show that such a partition is impossible. Let A be a part of $\mathcal{P}_{r_0+1}$

[1] The main conceptual difference is that we prove that each set A_i has no sparse cuts, whereas in [18] it is proved that each $G[A_i]$ is a robust expander; in fact, the latter would work for us as well, but we chose the former to simplify the presentation.

of the least order. Clearly, $|A| \le n/((1/c)+1) = cn/(c+1)$, and so every vertex in A has at least $cn(1-1/(c+1)) = c^2n/(c+1) \ge c_{\min}^2 n/2$ neighbours outside of A. Moreover, property (iv) implies that $|A| \ge 3^{-r_0} \cdot c_{\min} n$. Therefore, property (iii) implies that

$$3\sqrt{\eta_{r_0}} n^2 \;\ge\; |A| \cdot (c_{\min}^2\, n/2) \;\ge\; \frac{1}{2}\, 3^{-r_0}\, c_{\min}^2\, n^2,$$

contradicting the assumption that $\eta_{r_0} \ll c_{\min}$.

Consider the final partition $\mathcal{P}_r$. It consists of $r \le r_0$ parts, none of which have η_r-sparse cuts. We set $\zeta = \eta_r$, $\eta = 3\sqrt{\eta_{r-1}}$, $\delta = 3^{-r}c$ and observe that $\mathcal{P}_r$ satisfies properties (a) to (c) in Lemma 3. For property (d), we fix positive coefficients $\beta_0, \ldots, \beta_{r+1}$ that depend only on $c_{\min}$ and r, satisfying $3\sqrt{\eta_{r-1}} = \eta \ll \beta_0 \ll \cdots \ll \beta_{r+1} \ll \zeta = \eta_r$. For $i \in \{0, \ldots, r+1\}$, let $b(i)$ be the number of parts A in $\mathcal{P}_r$ that are β_i-almost-bipartite. Note that if A is β_i-almost-bipartite it is also β_{i+1}-almost-bipartite. In particular, $0 \le b(0) \le \ldots \le b(r+1) \le r$. It follows that there exists $i \in \{0, \ldots, r\}$ with $b(i) = b(i+1)$; fix one such i. Then every part A in $\mathcal{P}_r$ is either β_i-almost-bipartite or β_{i+1}-far-from-bipartite. Therefore, we can finish the proof by setting $\beta = \beta_i$ and $\gamma = \beta_{i+1}$. □

4 Hamiltonicity of clusters

Here the goal is to prove Lemma 4. Before turning to the proof, we mention the following proposition. A similar result can be deduced, for example, from Lemma 5.4 in [5]. We include a short proof, for completeness, in Appendix A.

Proposition 6. *Let $\zeta \in (0,1)$ and let H be a graph that has no ζ-sparse cuts. Then, for any $R \subset V(H)$ with $|R| \le (\zeta/6)|H|$ and any distinct vertices $x, y \in V(H) \setminus R$, there exists a path in $H \setminus R$ of length at most $3/\zeta$, with ends x and y.*

4.1 Preparing for the proof

We recall that n is the number of vertices of the fixed graph G. Given a graph H, a path P in H is called *ρ-absorbing* if, for every set of vertices $W \subset V(H) \setminus V(P)$ of size at most ρn, the induced subgraph $H[V(P) \cup W]$ contains a Hamilton path with the same ends as P. Similarly, if H is bipartite with bipartition $\{X, Y\}$, then P is called *ρ-bipartite-absorbing (with respect to $\{X, Y\}$)* if P has one end in X and one end in Y, and if the condition above holds for those $W \subset V(H) \setminus V(P)$ that satisfy $|(V(P) \cup W) \cap X| = |(V(P) \cup W) \cap Y|$ as well as $|W| \le \rho n$.

A result of DeBiasio and Nelsen [5] shows that absorbing paths exist in dense graphs of high edge-connectivity. We will use this result without giving its proof.

Lemma 7 (DeBiasio, Nelsen [5]). *Let $\rho, \gamma, \zeta, \delta$ be coefficients satisfying the hierarchy $1/n \ll \rho \ll \gamma \ll \zeta \ll \delta$ and let H be a graph with minimum degree at least δn and with no ζ-sparse cuts. Then H contains a path P of length at most ρn which is ρ^3-absorbing if H is γ-far-from-bipartite and ρ^3-bipartite-absorbing if H is bipartite.*

The following simple proposition enables us to join any pair of vertices with a short path, even if a moderate number of obstacles have to be avoided.

The following observation does not have much content, but it is a useful technical tool in situations where we throw away a small number of vertices and edges from a 'highly connected' graph. The observation says that, roughly speaking, the resulting graph is still highly connected.

Observation 8. *Let $\tau, \gamma, \delta, \zeta$ be real numbers satisfying $0 < \tau \ll \gamma \ll \zeta \ll \delta$ and let H be a graph on at most n vertices. Let H' be a graph obtained from H by removing at most τn^2 edges, and let $W \subset V(H)$ be a set of at most τn vertices.*

(i) *If H has no ζ-sparse cuts and the minimum degree of H' is at least δn, then the graph $H' \setminus W$ has no 0.99ζ-sparse cuts and its minimum degree is at least $0.99\delta n$.*

(ii) *If H is γ-far-from-bipartite, then $H' \setminus W$ is 0.99γ-far-from-bipartite. Conversely, if H is γ-almost-bipartite, then $H' \setminus W$ is γ-almost-bipartite.*

Proof. We first prove statement ((i)). If the minimum degree of H' is at least δn, then for any vertex $v \in V(H) \setminus W$ we have $d_{H' \setminus W}(v) \geq d_{H'}(v) - |W| \geq (\delta - \tau)n \geq 0.99\delta n$, as wanted. To prove the rest of the statement we assume that $\{A, B\}$ is a partition of $V(H) \setminus W$ such that $|A| \geq |B| \geq 1$ and $e(H'[A \cup B]) \leq 0.99\zeta |A||B|$. Since any vertex in B has at least $0.99\delta n - |B|$ neighbours in A, we have $|A| \geq 0.99\delta n - |B|$ and also $(0.99\delta n - |B|)|B| \leq 0.99\zeta |A||B| \leq 0.99\zeta n|B|$. It follows that $|B| \geq (0.99\delta - 0.99\zeta)n \geq 0.98\delta n$. Now, assuming that H has no ζ-sparse cuts, there are at least $\zeta|A \cup W||B|$ edges between $A \cup W$ and B in H. Combining this with the upper bound for $e(H'[A \cup B])$, we get $0.99\zeta|A||B| \geq \zeta|A \cup W||B| - \tau n^2 - n|W| \geq \zeta|A||B| - 2\tau n^2 \geq (\zeta - 2\tau/(0.98\delta)^2)|A||B|$, which contradicts the condition $\tau \ll \zeta \ll \delta$. This deals with statement ((i)).

It is clear that a subgraph of a γ-almost-bipartite graph is γ-almost-bipartite. Therefore, it remains only to prove the first part of statement ((ii)). Suppose that H is γ-far-from-bipartite and let $\{A, B\}$ be a partition of $V(H) \setminus W$. In H there are at least γn^2 edges with both ends in $A \cup W$ or in B, of which at most $\tau n^2 + n|W| \leq 2\tau n^2$ disappear when we pass to $H' \setminus W$. Therefore, in the latter graph there are at least $(\gamma - 2\tau)n^2 \geq 0.99\gamma n^2$ edges with both ends in A or in B, as desired. □

4.2 Proof of Lemma 4

We now prove Lemma 4. Our proof follows a strategy similar to one used by Lo and Patel [25] who showed that so-called robust expanders are Hamiltonian. In fact, Lemma 4 for γ-far-from-bipartite

A very nearly follows from Theorem 1.3 in [25], since under the conditions of the lemma, $G[A]$ is a robust expander; the only sense in which this case of our Lemma 4 is not fully covered by their Theorem 1.3 is that we can specify the ends of our Hamilton path. On the other hand, if A is β-almost-bipartite, then $G[A]$ is not a robust expander and the result of Lo and Patel is not applicable. Nevertheless, their strategy can be adapted to work for this case as well, and that is what we do in this subsection.

Proof of Lemma 4. Our setup depends on whether A is β-almost-bipartite or γ-far-from-bipartite. If A is β-almost-bipartite, then we fix a partition $\{X, Y\}$ of A that maximises the number of $X-Y$ edges and a set $W \subset A$ such that $|W| \leq \xi n$ and $|X \backslash W| = |Y \backslash W|$. We define $X_3 = X \backslash W,\ Y_3 = Y \backslash W$ and $H_3 = G[X_3, Y_3]$, and we pick vertices $x_* \in X_3,\ y_* \in Y_3$. (Admittedly, it is unusual to define X_3, Y_3, H_3 having not first introduced X_1, Y_1, H_1 or X_2, Y_2, H_2, but our notation makes it easier to keep track of various parameters: the minimum degree of H_i is at least $(\delta/i)n$ and H_i has no ζ/i-sparse cuts.) Since X, Y are chosen so that $e(G[X, Y])$ is maximised, the minimum degree of $G[X, Y]$ is at least $(\delta/2)n$. Since A has no ζ-sparse cuts and $e(G[A]) - e(G[X, Y]) \leq \beta n^2$, Observation 8 implies that H_3 has no $\zeta/3$-sparse cuts and that its minimum degree is at least $(\delta/3)n$. Obviously, H_3 is bipartite.

On the other hand, if A is γ-far-from-bipartite, then we define $H_3 = G[A \setminus W]$ and pick distinct vertices $x_*,\ y_* \in A \setminus W$. Observation 8 implies that H_3 is $\gamma/3$-far-from-bipartite and that, as in the previous case, H_3 has no $\zeta/3$-sparse cuts and its minimum degree is at least $(\delta/3)n$.

The definition of H_3 is illustrated in Figure 1.

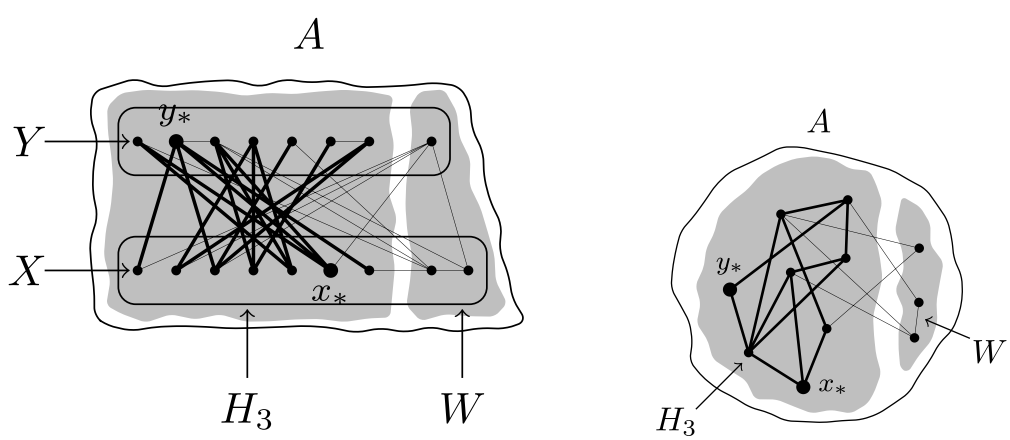

Figure 1: Definition of H_3. On the left – when A is β-almost-bipartite, on the right – when A is γ-far-from-bipartite.

Now, we fix a parameter ρ such that $1/n \ll \rho \ll \gamma$ and apply Lemma 7 to $H_3 \setminus \{x_*, y_*\}$. We get a path $Q \subset H_3 \setminus \{x_*, y_*\}$ of length at most ρn, where Q is ρ^3-absorbing if H_3 is $\gamma/3$-far-from-bipartite and Q is ρ^3-bipartite-absorbing with respect to the partition $\{\ X_3 \setminus \{x_*\},\ \ Y_3 \setminus \{y_*\}\ \}$ if H_3 is bipartite. Let x_0, y_0 be the ends of Q (with $x_0 \in X_3$, $y_0 \in Y_3$ if H_3 is bipartite) .

Now, we define $H_4 = H_3 \setminus (V(Q) \setminus \{x_0, y_0\})$ (see Figure 2) and, if H_3 is bipartite, $X_4 = (X_3 \setminus V(Q)) \cup \{x_0\}$, $Y_4 = (Y_3 \setminus V(Q)) \cup \{y_0\}$. By Observation 8, H_4 has no $\zeta/4$-sparse cuts and its minimum degree is at least $(\delta/4)n$; moreover, H_4 is either $\gamma/4$-far-from-bipartite or bipartite, depending on whether H_3 is $\gamma/3$-far-from-bipartite or bipartite.

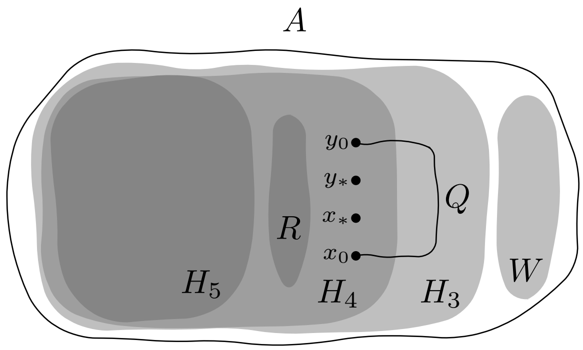

Figure 2: Definitions of H_3, H_4, H_5.

In order to find a Hamilton path in H_3 with ends x_*, y_*, we carry out the following three steps.

1. We set aside a small set of vertices $R \subset V(H_4) \setminus \{x_*, y_*, x_0, y_0\}$, called the *reservoir*, which can be used to join any given family of paths into a single path, provided that said family is not too large.

2. We find a cycle factor $\mathcal{F}$ for $H_4 \setminus (R \cup \{x_*, y_*, x_0, y_0\})$ that consists of not too many cycles. (A *cycle factor* for a graph is a collection of vertex-disjoint subgraphs, which are cycles, that cover all of its vertices.)

3. After removing an edge from each cycle, $\mathcal{F}$ becomes a family of paths. We use R to join Q and the paths in $\mathcal{F}$ into a single path $P \subset H_3$ with ends x_*, y_*. This path covers all of H_3 except for, possibly, some vertices in R. We absorb the uncovered part of R into Q (which, importantly, is a subpath of P). This step is illustrated in Figure 3.

Step 1.: *Choosing the reservoir R.*

This step can be completed via a rather straightforward probabilistic argument.

Claim 9. *There exists a set $R \subset V(H_4) \setminus \{x_0, y_0, x_*, y_*\}$ of size at most $(\log n)^3$ such that for every pair of distinct vertices $a, b \in V(H_4) \setminus R$ (with $a \in X_4$, $b \in Y_4$ if H_4 is bipartite) $H_4[R]$ contains $\log n$ vertex-disjoint paths, each of length at most $12/\zeta$, that start at neighbours of a and end at neighbours of b. Furthermore, if H_4 is bipartite, then we can also ensure that $|R \cap X_4| = |R \cap Y_4|$.*

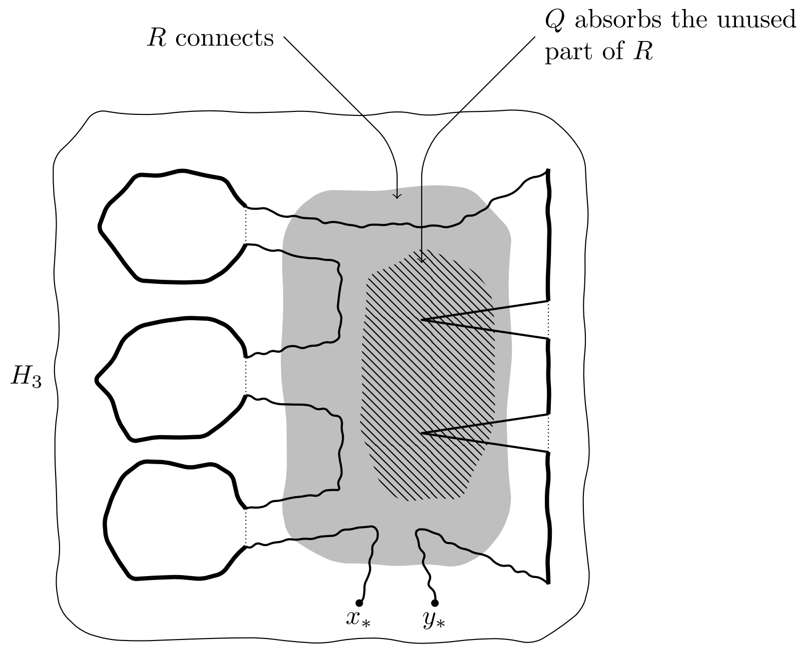

Figure 3: Construction of a Hamilton path in H_3 with ends x_*, y_*.

Proof. First, we choose a set $R_1 \subset V(H_4)$ at random by including each vertex of $V(H_4)$ independently with probability $(100/\delta)\log n/n$. Since $|H_4| \le n$, an application of a Chernoff's bound gives

$$\mathbb{P}\left[|R_1| \ge \frac{200}{\delta}\log n\right] \le \exp\left(-\frac{1}{3}\cdot\frac{100}{\delta}\log n\right) \le n^{-30/\delta} \ll 1.$$

Let v be an arbitrary vertex of H_4. We know that v has at least $(\delta/4)n - 4 \ge (24\delta/100)n$ neighbours in $H_4 \setminus \{x_0, y_0, x_*, y_*\}$, and hence

$$\mathbb{P}[v \text{ has at most } \log n \text{ neighbours in } R_1] \le \exp\left(-\frac{1}{8}\cdot 24\log n\right) = n^{-3}.$$

Taking the union bound over all $v \in V(H_3)$, we conclude that with high probability $|R_1| \le (200/\delta)\log n$ and every vertex of H_4 has at least $\log n$ neighbours in R_1. We fix an instance of R_1 with these properties. Furthermore, if H_4 is bipartite, then we make the two parts $R_1 \cap X_4$, $R_1 \cap Y_4$ have equal size by adding arbitrary vertices of H_4 to the smaller one; this results in R_1 satisfying the conditions $|R_1 \cap X_4| = |R_1 \cap Y_4|$ and $|R_1| \le (400/\delta)\log n$.

Now, for every pair of distinct vertices $u, v \in R_1$ (with $u \in X_4$, $v \in Y_4$ if H_4 is bipartite) we pick a path $P_{u,v} \subset H_4 \setminus \{x_0, y_0, x_*, y_*\}$ with ends u, v and of length at most $12/\zeta$ in such a way that the paths are pairwise internally vertex-disjoint (meaning that any two of them can intersect only at their ends). We pick these paths one by one by applying Proposition 6: when it is turn to select $P_{u,v}$ for a given pair $u, v \in V(R_1)$, we forbid vertices covered by the previously selected paths or contained in R_1, except for u, v. Proposition 6 applies, because H_4 has no $\zeta/4$-sparse cuts and we never forbid more than $(12/\zeta)|R_1|^2 + |R_1| + 4 \leq (\log n)^3$ vertices, while $|H_4| \geq (\delta/4)n$.

Finally, we take R to be the set of vertices covered by the paths selected in the previous paragraph. As in the previous calculation, $|R| \leq (\log n)^3$. Moreover, if H_4 is bipartite, then any path with ends in different parts of H_4 passes through the same number of vertices in both parts, and hence $|R \cap X_4| = |R \cap Y_4|$. $\square$

We fix R as given by Claim 9. We define $H_5 = H_4 \setminus (R \cup \{x_0, y_0, x_*, y_*\})$ (see Figure 2) and, if H_4 is bipartite, we also define $X_5 = X_4 \cap V(H_5)$, $Y_5 = Y_4 \cap V(H_5)$ (note that if X_5, Y_5 are defined, then $|X_5| = |Y_5|$). By Observation 8, H_5 has no $\zeta/5$-sparse cuts and its minimum degree is at least $(\delta/5)n$. Furthermore, H_5 is either $\gamma/5$-far-from-bipartite or bipartite with vertex classes X_5, Y_5. Also, recall that one of the assumptions of Lemma 4 is that G has at most ηn^2 edges with exactly one end in A. Since $V(H_5)$ was obtained by removing at most $\xi n + \rho n + 2 + (\log n)^3$ vertices from A, we conclude that G has at most $\eta n^2 + (\xi n + \rho n + 2 + (\log n)^3)|A| \leq 2\xi n^2$ edges with exactly one end in $V(H_5)$.

Step 2.: *Finding a cycle factor for H_5 with not too many cycles.*

The following claim states our aim precisely.

Claim 10. *There exists a cycle factor $\mathcal{F}$ for H_5 that consists of at most $3/\xi$ cycles.*

We shall achieve this bound by ensuring that every cycle in $\mathcal{F}$ covers at least $(\xi/3)n$ vertices. It turns out to be convenient to give direction to the cycles, and therefore we make the following definition. A *directed cycle factor* for H_5 is a directed graph on vertex set $V(H_5)$, possibly with loops and parallel arcs, satisfying

(a) every vertex has in-degree 1 and out-degree 1;

(b) upon erasing the directions of the arcs and removing loops and multiple edges we get a subgraph of H_5;

(c) if H_5 is bipartite, then there are no loops.

Conditions ((a)) and ((b)) are the important ones, but we also need ((c)) for technical reasons. Note that a directed cycle factor is a collection of vertex-disjoint directed cycles, some of which may be degenerate.

The following claim does all of the work in this step.

Claim 11. *Let $\mathcal{F}$ be a directed cycle factor for H_5 in which some component has cardinality at most $(\xi/3)n$. Then H_5 has a cycle factor with strictly fewer such components.*

Claim 10 quickly follows from a repeated application of this claim. However, before we show this deduction or even prove Claim 11, we introduce the concept of a robust neighbourhood, which will turn out to be important when dealing with ensuing technicalities.

Definition 12. Let $S \subseteq V(H_5)$. For each vertex $u \in S$, let $B(u)$ be a subset of $V(H_5)$, and set $\mathcal{B} = (B(u) : u \in S)$. We define the *$\mathcal{B}$-robust neighbourhood of S*, denoted by $N_{\mathcal{B}}(S)$, as

$$N_{\mathcal{B}}(S) = \bigcup_{u \in S} \left(N_{H_5}(u) \setminus B(u) \right).$$

Claim 13. *Let $S \subset V(H_5)$ be a non-empty set of size at most $|H_5| - (\delta/5)n$ if H_5 is $\gamma/5$-far-from-bipartite and at most $|H_5|/2 - (\delta/5)n$ if H_5 is bipartite. Let $\mathcal{B} = (B(u) : u \in S)$ be a collection of subsets of $V(H_5)$, each of size at most ξn. Then $|N_{\mathcal{B}}(S)| \geq |S| + \xi n$.*

Proof. We define $S_1 = S \setminus N_{\mathcal{B}}(S)$, $S_2 = S \cap N_{\mathcal{B}}(S)$, $N = N_{\mathcal{B}}(S) \setminus S$. Our first aim is to establish a lower bound on the number of edges between S_1 and N. Recall that in G there are at most $2\xi n^2$ edges between $V(H_5)$ and $V(G) \setminus V(H_5)$ and that H_5 either is an induced subgraph of G (if H_5 is $\gamma/5$-far-from-bipartite), or it was obtained by removing at most βn^2 edges from an induced subgraph of G (if H_5 is bipartite). Since G is cn-regular, we have $\sum_{u \in S_1} d_{H_5}(u) \geq cn|S_1| - 2\xi n^2 - 2\beta n^2 \geq cn|S_1| - 3\xi n^2$.

Throughout the rest of the proof we work in H_5. By the definition of the $\mathcal{B}$-robust neighbourhood, each $u \in S$ is adjacent to at most $|B(u)| \leq \xi n$ vertices in the complement of $N_{\mathcal{B}}(S)$. It follows that $e(S_1, N) \geq \sum_{u \in S_1}(d(u) - \xi n) - e(S_1, S_2)$ and also $e(S_1, S_2) \leq \xi n|S_2| \leq \xi n^2$. Therefore, $e(S_1, N) \geq cn|S_1| - 5\xi n^2$.

Let us assume for contradiction that $|N| < |S_1| + \xi n$. Since the maximum degree of H_5 is at most cn, the number of edges with one end in N and the other not in S_1 does not exceed $cn|N| - e(S_1, N) \leq 6\xi n^2$. Similarly, there are at most $cn|S_1| - e(S_1, N) \leq 5\xi n^2$ edges between S_1 and the complement of N, which means that there are at most $11\xi n^2$ edges joining $S_1 \cup N$ to the rest of H_5. Since H_5 has no $\zeta/5$-sparse cuts, we have the inequality $11\xi n^2 \geq (\zeta/5)|S_1 \cup N||V(H_5) \setminus (S_1 \cup N)|$. Recalling that $\xi \ll \zeta$, we conclude that the cardinality of $S_1 \cup N$ or its complement is at most $\xi^{1/3}n$.

There are at most $\xi n|S| \leq \xi n^2$ edges between S and the complement of $S \cup N$, because $S \cup N$ contains $N_{\mathcal{B}}(S)$. Moreover, a bound from the previous paragraph implies that there are at most $6\xi n^2$ edges between N and the complement of $S \cup N$. Therefore, $S \cup N$ can be separated from the rest of the graph by cutting at most $7\xi n^2$ edges. This gives the inequality $7\xi n^2 \geq (\zeta/5)|S \cup N||V(H_5) \setminus (S \cup N)|$. As S is non-empty, the minimum degree condition implies that $|S \cup N| \geq |N_{\mathcal{B}}(S)| \geq (\delta/5)n - \xi n \geq (\delta/6)n$, and hence the complement of $S \cup N$ contains at most $(210\ \xi/(\delta\zeta))n \leq \sqrt{\xi}n$ vertices. It follows that $2|S| + \xi n \geq |S| + |N| \geq |H_5| - \sqrt{\xi}n$ and, in particular, $|S| \geq |H_5|/2 - \sqrt{\xi}n$.

If H_5 is bipartite, then the last bound in the previous paragraph contradicts the assumption on the cardinality of S. Therefore, it remains to consider the case where H_5 is $\gamma/5$-far-from-bipartite.

We have $|N| \geq (|H_5| - \sqrt{\xi}n) - |S| \geq (\delta/5 - \sqrt{\xi})n > \xi^{1/3}n$, and hence the complement of $S_1 \cup N$ contains at most $\xi^{1/3}n$ vertices. It follows trivially that there are at most $\xi^{2/3}n^2$ edges not incident with $S_1 \cup N$. Moreover, we know that there are at most $11\xi n^2$ edges incident with S_1 but not with N or vice versa. Therefore, by removing at most $(\xi^{2/3} + 11\xi)n^2 < (\gamma/5)n^2$ edges we can leave only the $S_1 - N$ edges in H_5, thereby making it bipartite. However, this contradicts the assumption that H_5 if $\gamma/5$-far-from-bipartite. □

Proof of Claim 10 (assuming Claim 11). Provided that H_5 has a directed cycle factor, we repeatedly apply Claim 11 to it until there are no components of cardinality smaller than $(\xi/3)n$. In particular, the resulting directed cycle factor has no loops nor parallel arcs, and hence upon erasing the directions of the arcs we obtain a desired cycle factor for H_5.

It remains to construct a directed cycle factor for H_5. If H_5 is not bipartite, then this task is trivial: attaching a loop to each vertex of H_5 does the job. So suppose that H_5 is bipartite with vertex classes X_5, Y_5. We seek a perfect matching in H_5, as one can be converted to a directed cycle factor by replacing each edge with a pair of parallel arcs. Therefore, our task is reduced to checking Hall's condition for X_5. Let $S \subset X_5$ be a non-empty set. If $|S| \leq |X_5| - (\delta/5)n$, then by applying Claim 13 with all $B(u)$ empty we get $N(S) \geq |S|$, as wanted. On the other hand, if $|S| > |X_5| - (\delta/5)n$, then by the minimum degree condition every vertex in Y_5 is adjacent to something in S, and therefore $|N(S)| = |Y_5| = |X_5| \geq |S|$, also as wanted. □

The proof of Claim 11 is based on a fairly natural rotation procedure, which swaps an existing edge of a directed cycle factor for a new one. The following definition is useful. A *prefactor* (see Figure 4) is a directed graph on vertices $V(H_5)$, possibly with loops and parallel arcs, satisfying

(a) all vertices, except for one called the *root*, have in-degree 1; the root has in-degree 0;

(b) all vertices, except for one called the *pivot*, have out-degree 1; the pivot has out-degree 0;

(c) upon erasing the directions of the arcs and removing loops and multiple edges we get a subgraph of H_5;

(d) if H_5 is bipartite, then there are no loops.

Let $\mathcal{P}$ be a prefactor with root a and pivot b. A *rotation* of $\mathcal{P}$ is a procedure that involves picking a vertex $x \in V(H_5) \setminus \{a\}$ such that $bx \in E(H_5)$ and replacing the arc of $\mathcal{P}$ pointing to x with a new arc $\overrightarrow{bx}$. Note that the resulting prefactor has root a, but its pivot is not b (in fact, it is the in-neighbour of x in $\mathcal{P}$).

A rotation is *valid* if the selected vertex x is either not in the component of $\mathcal{P}$ containing a and b, or if x is in that component – which is a directed path starting at a and ending at b – but not one of its initial $(\xi/3)n$ nor terminal $(\xi/3)n$ vertices.

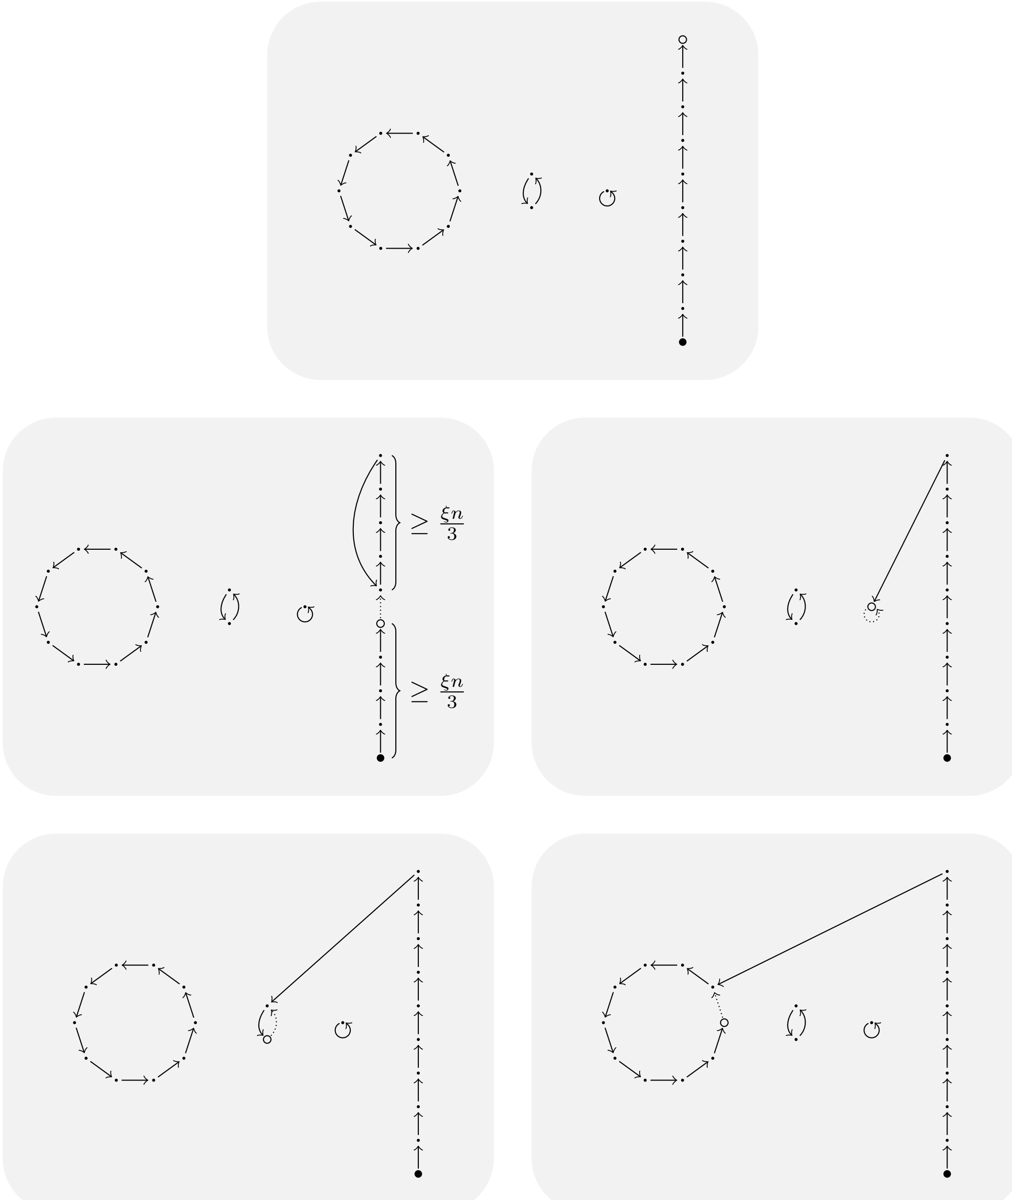

Figure 4: The top panel shows a prefactor $\mathcal{P}$. The other four panels show different valid rotations applied to $\mathcal{P}$. The root and the pivot are represented by a full dot and a hollow dot, respectively.

Proof of Claim 11. Fix a vertex $a \in V(H_5)$ in a component of $\mathcal{F}$ that has cardinality at most $(\xi/3)n$. Let $\mathcal{P}$ be the prefactor obtained by removing the arc of $\mathcal{F}$ that points to a. Note that a is the root of $\mathcal{P}$. For every $v \in V(H_5) \setminus \{a\}$, the *predecessor* of v is its in-neighbour in $\mathcal{P}$.

Our aim is to apply a sequence of valid rotations to $\mathcal{P}$ so that the resulting prefactor has pivot adjacent to a in H_5. Indeed, the addition of the arc directed from its pivot to a makes that prefactor into a directed cycle factor for H_5; moreover, provided that at least one valid rotation is applied, the number of components of cardinality at most $(\xi/3)n$ decreases.

To this end, for each $i \geq 0$ we say that a prefactor is *i-rotated* if it can be obtained by applying a sequence of i valid rotations to $\mathcal{P}$. Moreover, we define S_i as the set of vertices that appear as the pivot of an i-rotated prefactor.

We shall show that $S_i \cap N_{H_5}(a)$ is non-empty for some $1 \leq i \leq (\xi/3)n$. Let us suppose for contradiction that $S_i \cap N_{H_5}(a) = \emptyset$ for all $1 \leq i \leq (\xi/3)n$ and let us fix for the moment one value for i in this range. For each $u \in S_{i-1}$ we pick a witness $\mathcal{P}_u$ for u being in S_{i-1}; that is, $\mathcal{P}_u$ is an $(i-1)$-rotated prefactor with pivot u. In particular, $\mathcal{P}_u$ has at most $i-1$ arcs that are not present in $\mathcal{P}$. We define $B(u)$ as the set consisting of two types of vertices: (1) heads of the arcs present in $\mathcal{P}_u$, but not in $\mathcal{P}$, and (2) the initial $(\xi/3)n$ and terminal $(\xi/3)n$ vertices on the path component of $\mathcal{P}_u$ containing a, u. Clearly, $|B(u)| \leq \xi n$. Now, let $v \in N_{H_5}(u) \setminus B(u)$ and let w be the predecessor of v. Then, $\overrightarrow{wv}$ is an arc of $\mathcal{P}$ and also of $\mathcal{P}_u$. The rotation of $\mathcal{P}_u$ that replaces $\overrightarrow{wv}$ with $\overrightarrow{uv}$ is valid, and it produces an i-rotated prefactor with pivot w. Therefore, S_i contains the predecessor of each vertex in $N_{H_5}(u) \setminus B(u)$. Writing $\mathcal{B} = (B(u) : u \in S_{i-1})$, we see that, in fact, S_i contains the predecessors of all vertices in $N_{\mathcal{B}}(S_{i-1})$. Since no two vertices have the same predecessor, we have $|S_i| \geq |N_{\mathcal{B}}(S_{i-1})|$.

We are approaching a contradiction, but the way we reach it depends on whether H_5 is bipartite or $\gamma/5$-far-from-bipartite. Let us first consider the latter case, which is simpler. By the assumption that $S_i \cap N_{H_5}(a) = \emptyset$, we have $|S_i| \leq |H_5| - (\delta/5)n$ for all $1 \leq i \leq (\xi/3)n$. Claim 13 implies that $|S_i| \geq |S_{i-1}| + \xi n$, and so $|S_i| \geq \xi n i$ for all i in this range. However, this contradicts the obvious fact that $|S_i| \leq n$.

It remains to deal with the case when H_5 is bipartite. We may assume that $a \in X_5$. Let $\mathcal{P}'$ be a prefactor with root a and pivot b. By the definition of a prefactor, all arcs of $\mathcal{P}'$ are incident with both X_5 and Y_5. Since $|X_5| = |Y_5|$ and since each vertex except for a, b is incident with two arcs of $\mathcal{P}'$, the pivot b is in Y_5. Therefore, $S_i \subset Y_5$ for all $i \geq 0$. Now, we are done: since S_i does not intersect $N_{H_5}(a)$, we have $|S_i| \leq |Y_5| - (\delta/5)n$ for all $1 \leq i \leq (\xi/3)n$. Claim 13 implies that $|S_i| \geq |S_{i-1}| + \xi n$, and hence $|S_i| \geq \xi n i$ for all i in this range, giving a contradiction. □

Step 3.: *Connecting and absorbing (see Figure 3).*

By the previous step, there exists a cycle factor for H_5 that consists of at most $3/\xi$ cycles. By removing an edge from each of these cycles, we get a partition of $V(H_5)$ into paths $P_1, \ldots, P_k$ where $k \leq 3/\xi$. For each $1 \leq i \leq k$, we denote the ends of P_i by x_i, y_i, and we denote the ends of the absorbing path Q by x_0, y_0 (with $x_i \in X_5$ and $y_i \in Y_5$ for all i if H_5 is bipartite). Recall that x_* and y_* are the prescribed vertices of H_5, which we wish to connect by a Hamilton path.

By the choice of R, there exist paths $Q_0, \ldots, Q_{k+1}$ of length at most $12/\zeta$ each, whose interiors are pairwise disjoint subsets of R, such that

- Q_0 has ends x_0 and y_*;
- for each $1 \leq i \leq k$, Q_i has ends x_i and y_{i-1};
- Q_{k+1} has ends x_* and y_k.

More precisely, we pick the paths Q_i one by one: when we are about to pick Q_i, by the choice of R there are at least $\log n$ internally vertex-disjoint paths with the required ends, each of length at most $12/\zeta$, and at most $12i/\zeta < \log n$ of them intersect the previously chosen Q_j's.

Let P be the concatenation of paths $Q_0, Q, Q_1, P_1, \ldots, P_k, Q_{k+1}$ in this order. This path has ends x_*, y_*, is contained in $H_3 = H \setminus W$, and covers all vertices of H_3, except for, possibly, some vertices in R. To obtain the required Hamiltonian path, we absorb the vertices in R that are not covered by P into the path Q. If H_3 is not bipartite, then this is trivially possible since $|R| \leq (\log n)^3$ and Q is ρ^3-absorbing. If H_3 is bipartite, then we need to check that X_3 and Y_3 have the same number of uncovered vertices. However, this is clearly true since H_3 is bipartite and thus the path P visits an equal number of vertices in the two parts.

This concludes the proof of Lemma 4. □

We remark that vertices absorbed by the path Q are always vertices of R. Hence a more efficient way to prove Lemma 4 would be to first pick R, then pick an absorbing path for R, and then proceed as in the proof. Using this modification, it is possible to prove a suitable version of Lemma 4 for $d \geq n^{1-\varepsilon}$ for some constant $\varepsilon > 0$. However, in order to prove such a result, we would have to prove a variant of Lemma 7, rather than use it as a black box. As this would make the proof longer than it is now, we choose not to do so.

5 Balancing the bipartite clusters

In this section we prove Lemma 5. The proof spans the whole section and consists of several claims.

5.1 The setup

We first recap the setup needed for the proof of Lemma 5. We are given parameters $c_{\min}, n, \eta, \beta, \xi, \gamma, \zeta, \delta$ such that

$$1/n \ll \eta \ll \beta \ll \xi \ll \gamma \ll \zeta \ll \delta \ll c_{\min}.$$

We are also given a d-regular graph G, where $d \geq c_{\min} n$, and we denote $d = cn$, so that $c \geq c_{\min}$. We are further given a partition $\{A_1, \ldots, A_r\}$ of $V(G)$, where $r \leq \lceil 1/c_{\min} \rceil$, that satisfies the following properties.

(a) G has at most ηn^2 edges with ends in separate clusters;

(b) for each $i \in [r]$, the minimum degree of $G[A_i]$ is at least δn;

(c) for each $i \in [r]$, A_i has no ζ-sparse cuts;

(d) for each $i \in [r]$, either A_i is β-almost-bipartite, in which case we fix $\{X_i, Y_i\}$ to be a partition of A_i that maximises the number of $X_i - Y_i$ edges, or A_i is γ-far-from-bipartite.

For the sake of the proof of Lemma 16, ξ denotes any parameter satisfying $\beta \ll \xi \ll \gamma$; we will not use the fact that each A_i is ξ-Hamiltonian if it is γ-far-from-bipartite, or ξ-weakly-Hamiltonian if it is β-almost-bipartite, which follows from Lemma 4.

Our aim is to find a linear forest H in G, with the following properties.

(a) $|H| \leq \xi n$;

(b) H has no isolated vertices;

(c) for each $i \in [r]$, A_i contains either two or zero leaves of H;

(d) for each $i \in [r]$ such that A_i is β-almost-bipartite, either A_i contains no leaves of H, or X_i and Y_i each contain exactly one leaf of H;

(e) for each $i \in [r]$ such that A_i is β-almost-bipartite, $|X_i \setminus V(H)| = |Y_i \setminus V(H)|$.

In the proof, we shall consider the *lift* of G, denoted $\bar{G}$, which is a bipartite analogue of G. The lift $\bar{G}$ is defined as follows. We set $V(\bar{G}) = V^{(1)} \cup V^{(2)}$ where $V^{(1)}, V^{(2)}$ are disjoint copies of $V(G)$; for every $i \in \{1,2\}$ and $v \in V(G)$ we denote by $v^{(i)}$ the copy of v in $V^{(i)}$. For all $u, v \in V(G)$, $u^{(1)}v^{(2)}$ is an edge of $\bar{G}$ if and only if uv is an edge of G. There are no edges in $\bar{G}$ with both ends in $V^{(1)}$ or in $V^{(2)}$. It is clear from this construction that $\bar{G}$ is a cn-regular bipartite graph on $2n$ vertices.

The use of the lift $\bar{G}$ of G is convenient for us for three reasons. First, we shall be dealing with flows and matchings, and the fact that $\bar{G}$ is bipartite makes it easier to analyse them. Second, the lift allows us to treat β-almost-bipartite and γ-far-from-bipartite clusters in a unified way. And, third, consider a β-almost-bipartite cluster A_i, with prescribed partition $\{X_i, Y_i\}$, and suppose that $|Y_i| = |X_i| + k$. It turns out that in order to 'balance' A_i, it suffices to find two matchings M_1, M_2, whose union does not span any cycles or double edges, and, for $j \in \{1,2\}$, we have $|V(M_j) \cap Y| = |V(M_j) \cap X|$. Because $\bar{G}$ contains two copies of each such cluster, a so-called balancing matching for $\bar{G}$ (we make the notion precise below), pulled back to G, provides us with such 'overbalancing' automatically. In particular, if G is bipartite, then $\bar{G}$ consists of two copies of G, and its analysis allows us to find two such matchings simultaneously.

We partition the vertices of $\bar{G}$ into sets $\bar{A}_1, \dots, \bar{A}_s$, which we call *clumps* (which are related to, but should not to be confused with clusters $A_1, \dots, A_r$), as follows. Let i be the index of an arbitrary β-almost-bipartite cluster A_i of G and fix a partition $\{X_i, Y_i\}$ of A_i which maximises the number of $X_i - Y_i$ edges in G. In particular, $X_i, Y_i \neq \emptyset$ and all but at most βn^2 edges of $G[A_i]$ are between X_i and Y_i. Furthermore, every vertex of X_i (resp. Y_i) has at least $\delta n/2$ neighbours in Y_i (resp. X_i), as otherwise we could move that vertex to the other part, increasing the number of $X_i - Y_i$ edges. For $j \in \{1, 2\}$, let $X_i^{(j)}, Y_i^{(j)}$ be the copies of, respectively, X_i, Y_i in $V^{(j)}$. We define sets

$$\bar{A}_{i,1} = B_{i,1} \cup T_{i,1}, \text{ where } B_{i,1} = X_i^{(1)} \text{ and } T_{i,1} = Y_i^{(2)},$$
$$\bar{A}_{i,2} = B_{i,2} \cup T_{i,2}, \text{ where } B_{i,2} = Y_i^{(1)} \text{ and } T_{i,2} = X_i^{(2)}.$$

Now, let i be the index of some γ-far-from-bipartite cluster A_i. We define B_i and T_i to be the copies of A_i in $V^{(1)}$ and $V^{(2)}$, respectively, and

$$\bar{A}_i = B_i \cup T_i.$$

In these definitions B stands for the 'bottom part' and T stands for the 'top part'.

By doing this for all $i \in [r]$ we obtain a partition of $V(\bar{G})$ into clumps labelled $\bar{A}_{i,1}, \bar{A}_{i,2}$ (for those i for which A_i is β-almost-bipartite) and $\bar{A}_i$ (for the other i). To make the notation consistent, we relabel these clumps simply as $\bar{A}_1, \dots, \bar{A}_s$, where $s = r + |\{i \in [r] : A_i \text{ is } \beta\text{-almost-bipartite}\}|$. In particular, $s \in \{r, \dots, 2r\}$. We relabel the sets $B_{\dots}$ and $T_{\dots}$ appropriately, so that $\bar{A}_j = B_j \cup T_j$ for all $j \in [s]$.

Observation 14. *$\bar{G}$ has at most $3r\beta n^2$ edges with ends in separate clumps.*

Proof. First, note that every edge with both ends in a γ-far-from-bipartite cluster A_i of G gives rise to two edges of $\bar{G}$, both contained in the clump corresponding to A_i. Now, consider an arbitrary β-almost-bipartite cluster A_j of G. We recall that A_j is partitioned into sets X_j, Y_j such that all but at most βn^2 edges of $G[A_j]$ are $X_j - Y_j$ edges. In $\bar{G}$, A_j gives rise to two clumps, say, $\bar{A}_{j_1}$ and $\bar{A}_{j_2}$. If $e \in E(G[A_j])$ is an $X_j - Y_j$ edge, then e corresponds to two edges of $\bar{G}$, one in $\bar{A}_{j_1}$ and one in $\bar{A}_{j_2}$. Therefore, only those edges of $G[A_j]$ that are not $X_j - Y_j$ edges give rise to edges of $\bar{G}$ with ends in separate clumps. Also, we have to account for the edges of G that have ends in separate clusters. Thus, the number of edges of $\bar{G}$ with ends in separate clumps is at most $2\eta n^2 + 2r\beta n^2 \leq 3r\beta n^2$, using the assumption that $\eta \ll \beta$. $\square$

Observation 15. *For each $i \in [s]$, the minimum degree of $\bar{G}[\bar{A}_i]$ is at least $\delta n/2$. In particular, every vertex in $\bar{A}_i$ has at most $(c - \delta/2)n$ neighbours in $V(\bar{G}) \setminus \bar{A}_i$.*

Proof. Pick i and let A_j be the cluster of G that gives rise to $\bar{A}_i$. Let $v^{(t)}$ be an arbitrary vertex in $\bar{A}_i$, where $v \in A_j$, $t \in \{1,2\}$. If A_j is γ-far-from-bipartite, then v has at least δn neighbours in A_j, and every such neighbour u gives rise to the vertex $u^{(3-t)} \in \bar{A}_i$, which is adjacent to $v^{(t)}$.

So suppose that A_j is β-almost-bipartite with partition $A_j = X_j \cup Y_j$. We recall that this partition was chosen so that every vertex in X_j has at least $\delta n/2$ neighbours in Y_j and vice versa. Therefore, the number of $X_j - Y_j$ edges incident with v is at most $\delta n/2$, and, for every such edge uv, the vertex $u^{(3-t)}$ is a neighbour of $v^{(t)}$ in $\bar{G}[\bar{A}_i]$.

This proves the first part of the observation. Together with the fact that $\bar{G}$ is cn-regular, it implies the second part as well. $\square$

Let H be a bipartite graph with bipartition $\{X, Y\}$ and let $U \subset V(H)$. We define

$$\mathrm{imb}_H(U) = \big||U \cap X| - |U \cap Y|\big|.$$

We call this quantity the *imbalance* of U in H. If H is clear from the context, then we may write $\mathrm{imb}(U)$ instead of $\mathrm{imb}_H(U)$. Furthermore, we say that a subgraph $F \subset H$ *balances* U if $|(U \cap X) \setminus V(F)| = |(U \cap Y) \setminus V(F)|$.

To make sure that imbalance is well-defined, we adopt the convention that every bipartite graph comes with a prescribed vertex bipartition. This choice will usually be clear from the context. For example, $\bar{G}$ has bipartition $\{V^{(1)}, V^{(2)}\}$ and so does every relevant spanning subgraph of $\bar{G}$.

Now comes a key definition. Let σ be an ordering of $V(G)$. We define the spanning subgraph G_σ of $\bar{G}$ by setting

$$E(G_\sigma) = \Big\{u^{(1)}v^{(2)} \,:\, uv \in E(G),\ \sigma(u) < \sigma(v) \text{ and } u^{(1)}, v^{(2)} \text{ are in distinct clumps of } \bar{G}\Big\}.$$

The rest of the proof goes as follows. First, we show that there exists an ordering σ of $V(G)$ such that G_σ contains a so-called balancing matching (see Lemma 16). The reason we consider G_σ instead of working directly with $\bar{G}$ is that a matching in G_σ of size m corresponds to a linear forest in G of size m, whereas the edges of G corresponding to a matching in $\bar{G}$ may span a cycle; moreover, an edge uv in G may be represented twice in a matching in $\bar{G}$ – once as $u^{(1)}v^{(2)}$ and once as $u^{(2)}v^{(1)}$. We explain this more precisely towards the end of the section. Second, we take the linear forest in G that comes from a balancing matching in G_σ, and we modify it slightly so that it satisfies the assertions of Lemma 5.

5.2 Balancing G_σ

Here comes the main technical lemma of the section.

Lemma 16. *There is an ordering σ such that G_σ has a matching M with the following properties:*

(a) *for each $i \in [s]$, M balances $\bar{A}_i$;*

(b) $|M| \leq (\xi\zeta/8)n$.

Property (a) is the main part of this lemma: if we find a matching in G_σ that balances $\bar{A}_1, \ldots, \bar{A}_s$, then we get property (b) for free from the following argument.

Proposition 17. *Let H be a balanced bipartite graph whose vertex set is partitioned into sets $U_1, \ldots, U_k$. Suppose that M is a matching in H that balances U_i for every $i \in [k]$. Then M contains a matching that has at most $(k-1)(\mathrm{imb}(U_1) + \cdots + \mathrm{imb}(U_k))$ edges and balances U_i for every $i \in [k]$.*

Proof. We use induction on k. The base case is when $k = 1$, in which case $U_1 = V(H)$ and $\mathrm{imb}(U_1) = 0$ because H is balanced, so the empty matching is a balancing matching.

Next, suppose that $k \geq 2$. Denote the bipartition of H by $\{X, Y\}$, and let $X_i = X \cap U_i$, $Y_i = Y \cap U_i$ for $i \in [k]$. Without loss of generality, we assume that M is a minimal matching that balances U_i for every $i \in [k]$. We claim that there is $i \in [k]$ for which M does not touch Y_i. Indeed, let D be an auxiliary directed graph on vertex set $[k]$, where ij is an edge if there is an X_i–Y_j edge in M. Suppose that $(i_1 \ldots i_s)$ is a directed cycle in D. Then there exist $a_j \in X_{i_j}, b_j \in Y_{i_j}$ such that $a_1b_2, \ldots, a_sb_1$ are edges of M. But then $M \setminus \{a_1b_2, \ldots, a_sb_1\}$ balances every U_i, contradicting the minimality of M. It follows that D is acyclic, which implies the existence of $i \in [k]$ with in-degree 0 in D, i.e. M does not touch Y_i, as claimed.

Without loss of generality, suppose that M does not touch Y_k. As M balances U_k, the number of edges of M that touch X_k is exactly $|X_k| - |Y_k| = \mathrm{imb}(U_k)$. Let M' be the submatching of M obtained by removing the edges that touch X_k, let H' be the subgraph of H obtained by removing U_k and vertices of $U_1 \cup \cdots \cup U_{k-1}$ that are neighbours of X_k in M, and let $U'_i = U_i \cap V(H')$. Then H' is a balanced bipartite graph, as exactly $|X_k|$ vertices are removed from each part of H to form H'. Moreover, M' is a minimal matching in H' that balances U'_i for every $i \in [k-1]$. Thus, by induction,

$$|M'| \leq (k-2)(\mathrm{imb}(U'_1) + \cdots + \mathrm{imb}(U'_{k-1})) \leq (k-2)(\mathrm{imb}(U_1) + \ldots + \mathrm{imb}(U_k)),$$

because the sum of imbalances of $U_1, \ldots, U_{l-1}$ increases by at most $\mathrm{imb}(U_k)$ when going from H to H'. Since $|M \setminus M'| = \mathrm{imb}(U_k)$, we have $|M| \leq (k-1)(\mathrm{imb}(U_1) + \cdots + \mathrm{imb}(U_k))$, as required. □

Observation 18. $\sum_{i=1}^{s} \mathrm{imb}_{\bar{G}}(\bar{A}_i) \leq (6\beta r/c)n$.

Proof. Pick $i \in [s]$ and recall that T_i, B_i are the vertex classes of $\bar{A}_i$. Since $\bar{G}$ is cn-regular, we have

$$|T_i|cn = e(T_i, B_i) + e(T_i, V(\bar{G}) \setminus \bar{A}_i) \leq |B_i|cn + e(T_i, V(\bar{G}) \setminus \bar{A}_i)$$

From this upper bound for $|T_i|cn$ and the corresponding upper bound for $|B_i|cn$ we get

$$\mathrm{imb}_{\bar{G}}(\bar{A}_i) = \big||T_i| - |B_i|\big| \leq \frac{e(\bar{A}_i, V(\bar{G}) \setminus \bar{A}_i)}{cn}$$

Summing over all i and applying Observation 14 gives the desired result. □

Proof of Lemma 16. As noted above, it is enough to find an ordering σ and a matching $M \subset G_\sigma$ that satisfies property (a) in Lemma 16. Indeed, Proposition 17 and Observation 18 then give us a submatching of M that satisfies property (a) and has at most $(12\beta r^2/c)n \leq (\xi\zeta/8)n$ edges, the latter bound being a consequence of the assumption that $\beta \ll \xi \ll \zeta$. We split our proof into two main steps. In the first step we find an ordering σ for which there is an almost balancing fractional matching in G_σ. In the second step we convert it to a balancing matching in G_σ.

Step 1: *Using the Max-Flow Min-Cut theorem to obtain an almost balancing fractional matching in G_σ for some ordering σ.*

The terms used in the summary of this step are mostly self-explanatory, but we define them formally to clarify the details. A *fractional matching* in G_σ is a function w that assigns weights from the interval $[0, 1]$ to the edges of G_σ in such a way that for each vertex $v \in V(G_\sigma)$ the *weight* of v, denoted $w(v)$ and defined as $\sum_{uv \in E(G_\sigma)} w(uv)$, does not exceed 1. Let w be a fractional matching in G_σ. For any $U \subset V(G_\sigma)$ we define $w(U) = \sum_{v \in U} w(v)$. For each $i \in [s]$ we define $\mathrm{imb}(w, i) = \big|(|T_i| - w(T_i)) - (|B_i| - w(B_i))\big|$. We say that w is *α-balancing* if $\sum_{i=1}^{s} \mathrm{imb}(w, i) \leq \alpha$. One can think of $w(T_i)$ as the weight of the edges leaving T_i (recall that in G_σ there are no $T_i - B_i$ edges), and similarly for B_i. If $\mathrm{imb}(w, i) = 0$, this means that the fractional matching w balances the cluster U_i. Since we are not able to find such a balancing fractional matching directly, we settle for one that is nearly-balancing, and the quantity $\mathrm{imb}(w, i)$ allows us to measure how far w is from balancing U_i. In this step we will find a 0.9-balancing fractional matching in G_σ for some σ.

We now prepare G_σ for an application of the Max-Flow Min-Cut theorem, that is, we convert it to a weighted digraph $\vec{G}_\sigma$ with a source and a sink (see Figure 5). The vertex set of $\vec{G}_\sigma$ contains $V(\bar{G})$ and $2s + 2$ new vertices: source p, sink q and, for each $i \in [s]$, a pair of new vertices b_i, t_i. The edges of G_σ become arcs of $\vec{G}_\sigma$, directed from $V^{(1)}$ to $V^{(2)}$ (we recall that $V^{(1)} = B_1 \cup \cdots \cup B_s$ and $V^{(2)} = T_1 \cup \cdots \cup T_s$). For every $i \in [s]$ we add arcs (1) from p to b_i, (2) from b_i to all vertices in B_i, (3) from all vertices in T_i to t_i and (4) from t_i to q. Vertices that were present in G_σ get capacity 1, while p, q get infinite capacity. The capacities of b_i, t_i, $i \in [s]$, are defined via quantities

a_{ij}, $i, j \in [s]$, which we now introduce. We set

$$a_{ij} = \begin{cases} \frac{1}{cn}\, e_{\bar{G}}(B_i, T_j) & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

and, for every $k \in [s]$,

$$b_k \text{ gets capacity } \sum_{j=1}^{s} a_{kj}, \qquad\qquad t_k \text{ gets capacity } \sum_{i=1}^{s} a_{ik}.$$

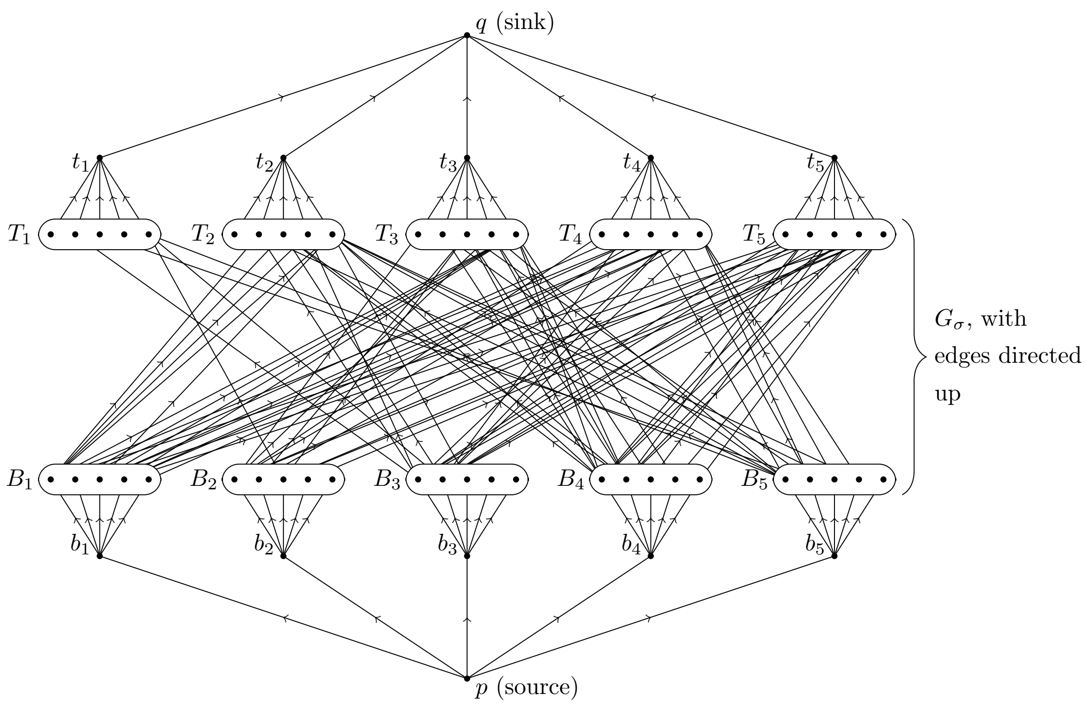

Figure 5: Definition of $\vec{G}_\sigma$.

A *cut* of $\vec{G}_\sigma$ is a subset of $V(\vec{G}_\sigma) \setminus \{p, q\}$ whose removal from $\vec{G}_\sigma$ disconnects q from p.

We will show that, for some σ, $\vec{G}_\sigma$ does not have cuts with capacity less than $\sum_i \sum_j a_{ij} - 0.9$. We will then apply the Max-Flow Min-Cut theorem to deduce the existence of a flow of at least this value, which in turn implies the existence of the required 0.9-balancing fractional matching. We note that the standard version of the Max-Flow Min-Cut theorem places capacities on the arcs than on the vertices and uses an appropriate notion of a cut. As the version that we use can be proved similarly to the standard one, we elect to use it out of convenience.

The choice of capacities of the vertices b_i and t_i may seem arbitrary at first glace[2], so before

[2]and, indeed, some trial and error was required in order to arrive at this 'correct' choice of capacities,

proceeding let us briefly explain why this choice makes sense. For each $i \in [s]$, the difference between the capacity of b_i and the capacity of t_i, is the difference between the number of edges of $\bar{G}$ incident with B_i and the number of edges incident with T_i, divided by cn, which is exactly the imbalance of the clump $\bar{A}_i$ in $\bar{G}$. It follows that a flow in G_σ that fully saturates both b_i and t_i (namely, the amount of flow through each of these vertices equals their capacity) translates into a fractional matching in $\bar{G}$ that balances $\bar{A}_i$. Thus, a flow in G_σ in which b_i and t_i are fully saturated for all $i \in [s]$, translates into the desired balancing fractional matching. Since the value of such a flow is $\sum_i \sum_j a_{ij}$, any flow with almost this value (a proof of whose existence in some G_σ will be the main aim of this section) almost saturates the vertices b_i and t_i for each $i \in [s]$, and translates into the required almost-balancing fractional matching.

With the goal of proving that some G_σ has no cuts of low capacity in mind, we consider graphs $F_{I,J,\sigma}$, defined for all $I, J \subset [s]$, that are the induced subgraphs of G_σ on vertices

$$V(F_{I,J,\sigma}) = \left(\bigcup_{i \in I} B_i \right) \cup \left(\bigcup_{j \in J} T_j \right).$$

The point of this definition is that every cut of $\vec{G}_\sigma$ induces a vertex cover of $F_{I,J,\sigma}$ for appropriately chosen I, J. This is why the following claim is useful.

Claim 19. *Fix $I, J \subset [s]$. Let σ be a random ordering of $V(G)$, chosen uniformly at random. With probability greater than $1 - 4^{-s}$, every vertex cover of $F_{I,J,\sigma}$ contains at least $\sum_{i\in I} \sum_{j \in J} a_{ij} - 0.9$ vertices.*

Proof. We define

$$E_{I,J} = \left\{ u^{(1)}v^{(2)} : uv \in E(G) \text{ and } u^{(1)} \in B_i, v^{(2)} \in T_j \text{ with } i \in I, j \in J, i \neq j \right\}.$$

In other words, $E_{I,J}$ is the set of edges of $\bar{G}[V(F_{I,J,\sigma})]$ that have ends in separate clumps. Note that $|E_{I,J}| = cn \sum_{i \in I} \sum_{j \in J} a_{ij}$ and that any given edge $u^{(1)}v^{(2)} \in E_{I,J}$ is in $F_{I,J,\sigma}$ if and only if $\sigma(u) < \sigma(v)$. Furthermore, it follows from Observation 15 that any vertex in $V(F_{I,J,\sigma})$ is incident with at most $(c - \delta/2)n$ edges in $E_{I,J}$.

We classify the vertices of $F_{I,J,\sigma}$ as rich or poor, according to the following rule (which does not depend on σ):

$$v \in V(F_{I,J,\sigma}) \text{ is } \begin{cases} \textit{rich} \text{ if } v \text{ is incident with at least } cn/(1000s) \text{ edges in } E_{I,J} \\ \textit{poor} \text{ otherwise.} \end{cases}$$

We also say that $e \in E_{I,J}$ is *rich* if at least one end of e is rich and *poor* otherwise. We write E_{rich} and E_{poor} to denote the sets of, respectively, rich and poor edges in $E_{I,J}$.

Our strategy is as follows: first, with high probability, we construct a matching in $E_{\text{poor}} \cap E(F_{I,J,\sigma})$ of size at least $|E_{\text{poor}}|/(cn) - 0.9$; then, also with high probability, we construct a matching in $E_{\text{rich}} \cap E(F_{I,J,\sigma})$ of size at least $|E_{\text{rich}}|/(cn)$, ensuring that these two matchings are vertex-disjoint. If we are successful in both tasks, then the union of these matchings is a matching in $F_{I,J,\sigma}$ of size at least $|E_{I,J}|/(cn) - 0.9$, giving the desired result.

First, we deal with the poor edges. Since the smallest vertex cover of E_{poor} contains only poor vertices, the cardinality of such a cover is at least $1000s|E_{\text{poor}}|/(cn)$. By Kőnig's theorem E_{poor} contains a matching M of size $|M| \geq 1000s|E_{\text{poor}}|/(cn)$. We say that two distinct edges $e, f \in M$ are *related* if there exists a vertex $u \in V(G)$ such that $u^{(1)}$ is an end of e and $u^{(2)}$ is an end of f, or vice versa. We greedily construct a subset $M' \subset M$ such that $|M'| \geq |M|/3$ and M' does not contain any pairs of related edges: initially we set $M' = \emptyset$ and consider the edges in M one by one, putting $e \in M$ into M' if e is not related to any edges already present in M'. The bound $|M'| \geq |M|/3$ comes from the fact that any edge of M is related to at most two other edges. Indeed, for every edge $e \in M \setminus M'$ there exists an edge in M' that prevented e from being accepted into M', while a single edge in M' can prevent at most two edges from being accepted, giving $|M \setminus M'| \leq 2|M'|$.

Let $\mathcal{E}_1$ be the event that $|M' \cap E(F_{I,J,\sigma})| \geq |E_{\text{poor}}|/(cn) - 0.9$. A given edge $u^{(1)}v^{(2)} \in M'$ is in $E(F_{I,J,\sigma})$ if and only if $\sigma(u) < \sigma(v)$, which happens with probability $1/2$. Moreover, since M' does not contain related edges, the events of particular edges of M' being present in $E(F_{I,J,\sigma})$ are independent, because they are determined by restrictions of σ to mutually disjoint pairs of vertices. As a result, $|M' \cap E(F_{I,J,\sigma})|$ has distribution $\text{Binom}(|M'|, 1/2)$. An application of a Chernoff's bound gives

$$\mathbb{P}\left[\left|M' \cap E(F_{I,J,\sigma})\right| < \frac{|M'|}{3}\right] \leq \exp\left(-\frac{|M'|}{18}\right).$$

Note that, in particular, $|M'|/3 \geq 1000s|E_{\text{poor}}|/(9cn) \geq |E_{\text{poor}}|/(cn)$. If $|E_{\text{poor}}| \geq (162/1000)cn$, then we also have $|M'| \geq 54s$, and hence $\mathcal{E}_1$ holds with probability at least $1 - \exp(-3s) > 1 - 4^{-s}/2$. On the other hand, if $|E_{\text{poor}}| < (162/1000)cn$, then $|E_{\text{poor}}|/(cn) < 0.9$, which means that $\mathcal{E}_1$ trivially holds. In either case,

$$\mathbb{P}(\mathcal{E}_1) > 1 - \frac{4^{-s}}{2}.$$

We now turn our focus to the rich edges. First, suppose that $E_{\text{rich}} \neq \emptyset$. Since any vertex in $V(F_{I,J,\sigma})$ is incident with at most $(c - \delta/2)n$ edges in $E_{I,J}$, there are at least $|E_{\text{rich}}|/(cn - \delta n/2)$ rich vertices. Let $\ell = \lceil |E_{\text{rich}}|/(cn - \delta n/2) \rceil$ and let R be a set of ℓ rich vertices. We say that a vertex in R is *ruined* if its degree in $F_{I,J,\sigma}$ is smaller than $\delta\sqrt{\beta}n$. Consider an arbitrary vertex in R that belongs to the vertex class $V^{(1)}$, that is, a vertex of the form $v^{(1)} \in R$ with $v \in V(G)$. Let $u_1, \ldots, u_d$ be the vertices in $V(G)$ such that $u_1^{(2)}, \ldots, u_d^{(2)}$ are adjacent to $v^{(1)}$ via edges in $E_{I,J}$. Since $v^{(1)}$ is rich, $d \geq cn/(1000s)$. Note that $v^{(1)}$ is ruined if and only if v appears in one of the final $\lceil \delta\sqrt{\beta}n \rceil$ positions of the order that σ induces on $\{v, u_1, \ldots, u_d\}$. Since v is equally likely to be in any position of this order, we have

$$\mathbb{P}\Big[v^{(1)} \text{ is ruined}\Big] \le \frac{\delta\sqrt{\beta}n+1}{cn/(1000s)+1} \le \frac{2000s\delta\sqrt{\beta}}{c} < \frac{\delta}{4^{s+1}c},$$

where the latter inequality comes from the assumption that $\beta \ll c_{\min}$. The same bound holds for those vertices in R that are in the vertex class $V^{(2)}$. Hence, the expected number of ruined vertices in R is at most $4^{-s-1}\delta\ell/c$. Markov's inequality gives

$$\mathbb{P}\Big[R \text{ has at least } \frac{\delta}{2c}\ell \text{ ruined vertices}\Big] < \frac{4^{-s}}{2}.$$

Let $\mathcal{E}_2$ be the event that at least $|E_{\text{rich}}|/(cn)$ vertices in R are not ruined. If $E_{\text{rich}} = \emptyset$, then $\mathcal{E}_2$ trivially holds. Otherwise, as we have just seen, with probability greater than $1 - 4^{-s}/2$, there are at least $(1 - \delta/(2c))\ell$ vertices in R that are not ruined. Since $\ell \ge |E_{\text{rich}}|/(cn - \delta n/2)$, we have $\mathbb{P}(\mathcal{E}_2) > 1 - 4^{-s}/2$.

At this point we have established that $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) > 1 - 4^{-s}$. We will finish the proof of the claim by assuming that $\mathcal{E}_1, \mathcal{E}_2$ both occur and constructing a matching in $F_{I,J,\sigma}$ of size at least $|E_{\text{poor}}|/(cn) + |E_{\text{rich}}|/(cn) - 0.9$. From $\mathcal{E}_1$ we get a matching $M_0 \subset E_{\text{poor}} \cap E(F_{I,J,\sigma})$ of size $|M_0| \ge |E_{\text{poor}}|/(cn) - 0.9$. Furthermore, since $\mathcal{E}_2$ occurs, there exist $m = \lceil |E_{\text{rich}}|/(cn) \rceil$ distinct rich vertices $v_1, \dots, v_m \in V(F_{I,J,\sigma})$ of degree at least $\delta\sqrt{\beta}n$ in $F_{I,J,\sigma}$. Note that $v_1, \dots, v_m \notin V(M_0)$ because the edges in M_0 are poor.

We now construct an eventually terminating sequence of matchings $M_0 \subset M_1 \subset \cdots$ in $F_{I,J,\sigma}$, where M_{i+1} is obtained by adding to M_i a single edge incident with v_{i+1}. Suppose that we have just constructed M_i for some $i \ge 0$. If $|M_i| \ge |E_{\text{poor}}|/(cn) + |E_{\text{rich}}|/(cn) - 0.9$, then we stop. If not, then we have $i \le m-1$, because $|M_i| = |M_0| + i$. Since v_{i+1} has at least $\delta\sqrt{\beta}n$ neighbours in $F_{I,J,\sigma}$, we can pick one, say u_{i+1}, that is not contained in $V(M_i) \cup \{v_{i+2}, \dots, v_m\}$ (here we use the bound $|V(M_i)| + m \le 3|E_{I,J}|/(cn) + 1 \le (9r\beta/c)n + 1 < \delta\sqrt{\beta}n$, which is a consequence of Observation 14 and the assumption that $\beta \ll \delta$). The new matching M_{i+1} is defined as $M_i \cup \{v_{i+1}u_{i+1}\}$. We remark that our construction ensures that at each stage v_{i+1} is not contained in $V(M_i)$, and so the process keeps running until we obtain a matching of a desired size. Claim 19 is proved. □

Claim 20. *There exists σ for which the capacity of every cut of $\vec{G}_\sigma$ is at least $\sum_{i=1}^{s}\sum_{j=1}^{s} a_{ij} - 0.9$.*

Proof. Let σ be a random ordering, chosen uniformly at random. For any $I, J \subset [s]$, let $\mathcal{E}_{I,J}$ be the event that $F_{I,J,\sigma}$ has no vertex cover of cardinality less than $\sum_{i\in I}\sum_{j\in J} a_{ij} - 0.9$. We know from the previous claim that $\mathbb{P}(\mathcal{E}_{I,J}) > 1 - 4^{-s}$ for any I, J. Since there are 4^s choices for I, J, all events $\mathcal{E}_{I,J}$ occur simultaneously with positive probability.

Suppose that $\mathcal{E}_{I,J}$ occurs for every $I, J \subset [s]$ and let C be a cut of $\vec{G}_\sigma$. Then $\mathcal{E}_{I,J}$ holds in particular for the choice $I = \{i \in [s] : b_i \notin C\}$, $J = \{j \in [s] : t_j \notin C\}$. Since C disconnects q from p, it in particular intersects all paths from p to q that visit $(\bigcup_{i\in I} B_i) \cup (\bigcup_{j\in J} T_i) = V(F_{I,J,\sigma})$, and hence $C \cap V(F_{I,J,\sigma})$ is a vertex cover of $F_{I,J,\sigma}$. Therefore,

$$\begin{aligned}
\text{capacity}(C) &= \sum_{i\notin I} \text{capacity}(b_i) + \sum_{j\notin J} \text{capacity}(t_j) + |C \cap V(F_{I,J,\sigma})| \\
&\geq \sum_{i\notin I}\sum_{j} a_{ij} + \sum_{i}\sum_{j\notin J} a_{ij} + \sum_{i\in I}\sum_{j\in J} a_{ij} - 0.9 \\
&\geq \sum_{i}\sum_{j} a_{ij} - 0.9,
\end{aligned}$$

where the first inequality follows from the assumption that $\mathcal{E}_{I,J}$ occurs. $\square$

We fix one instance of σ for which the capacity of a minimum cut of $\vec{G}_\sigma$ is at least $\sum_i \sum_j a_{ij} - 0.9$. The Max-Flow Min-Cut Theorem produces a flow f on $\vec{G}_\sigma$ with $\text{value}(f) \geq \sum_i \sum_j a_{ij} - 0.9$. This flow induces a fractional matching in G_σ, as the capacity of vertices in G_σ was set to 1. Abusing the notation slightly, we denote this fractional matching also by f.

Claim 21. *The fractional matching f is 0.9-balancing.*

Proof. It is clear from the way the directed graph $\vec{G}_\sigma$ was set up that, for every $i \in [r]$, $f(B_i)$ does not exceed the capacity of b_i in $\vec{G}_\sigma$. That is, $\sum_j a_{ij} - f(B_i) \geq 0$. Therefore,

$$\sum_i \sum_j a_{ij} - 0.9 \;\leq\; \text{value}(f) \;=\; f(B_1) + \cdots + f(B_r) \;\leq\; \sum_i \sum_j a_{ij},$$

from which we deduce that

$$0 \leq \sum_i \left(\sum_j a_{ij} - f(B_i) \right) \leq 0.9.$$

Similarly, we have $\sum_i a_{ij} - f(T_j) \geq 0$ for all j and

$$0 \leq \sum_j \left(\sum_i a_{ij} - f(T_j) \right) \leq 0.9.$$

At this point it is important to remember that for all distinct i, j we have $a_{ij}\, cn = e_{\bar{G}}(B_i, T_j)$. Also, $a_{ii} = 0$. Since $\bar{G}$ is cn-regular, for each $k \in [s]$ we have $|B_k|cn - |T_k|cn = \sum_j a_{kj}\, cn - \sum_i a_{ik}\, cn$, which can be rearranged to give

$$\begin{aligned}
\text{imb}(f,k) &= \left| \left(\sum_j a_{kj} - f(B_k) \right) - \left(\sum_i a_{ik} - f(T_k) \right) \right| \\
&\leq \max \left\{ \sum_j a_{kj} - f(B_k),\; \sum_i a_{ik} - f(T_k) \right\}.
\end{aligned}$$

Therefore,

$$\sum_k \operatorname{imb}(f,k) \le \max\left\{\sum_k \left(\sum_j a_{kj} - f(B_k)\right), \sum_k \left(\sum_i a_{ik} - f(T_k)\right)\right\} \le 0.9,$$

as claimed. $\square$

Step 2: *Converting the almost balancing fractional matching to a balancing matching.*

Let w be any fractional matching in G_σ. We say that a vertex $v \in V(G_\sigma)$ is *open* if $w(v) \in (0,1)$ and *closed* if $w(v) \in \{0,1\}$. Similarly, we say that an edge $e \in E(G_\sigma)$ is *open* if $w(e) \in (0,1)$ and *closed* if $w(e) \in \{0,1\}$.

We know that G_σ has a 0.9-balancing fractional matching, namely, f. Let f^* be a 0.9-balancing fractional matching in G_σ that minimises the total number of open vertices and open edges.

Claim 22. *The fractional matching f^* is 0-balancing and has integer weights.*

Proof. It suffices to show that f^* has no open edges. Indeed, this would imply that f^* has no open vertices, and so its imbalance is a whole number. However, by definition, $\operatorname{imb}(f^*) \le 0.9$, and so $\operatorname{imb}(f^*) = 0$, as required.

We assume for contradiction that f^* has at least one open edge. Let E_{open} and V_{open} stand for the sets of, respectively, open edges and open vertices of f^*. There may be closed vertices that are incident with open edges; we call such vertices *full* and denote their set by V_{full}. Clearly, full vertices have weight 1 and are incident with at least two open edges. We will now add new edges, which we call *fake*, to G_σ. For every $i \in [s]$, we add a path spanning the open vertices contained in $\bar{A}_i$. In particular, if for some i there is at most one open vertex in $\bar{A}_i$, then we do not create any fake edges in the clump $\bar{A}_i$. Let E_{fake} stand for the set of fake edges that were added to G_σ.

We create an auxiliary graph H with vertices $V_{\text{open}} \cup V_{\text{full}}$ and edges $E_{\text{open}} \cup E_{\text{fake}}$. First, suppose that H contains a cycle C. Since E_{fake} is a union of vertex-disjoint paths, C must contain at least one open edge. Fix a direction for C. For e an open edge in C, we say that e is *upward* if it is directed from $V^{(1)}$ to $V^{(2)}$. If e is directed from $V^{(2)}$ to $V^{(1)}$, then we say that e is *downward*.

Let $\lambda > 0$ be a small positive number and let f^*_λ be the fractional matching in G_σ obtained from f^* by adding λ to the weight of every upward open edge and subtracting λ from the weight of every downward open edge. We remark that f^*_λ is a valid fractional matching, provided that λ is small enough so that the modified weights of open edges and open vertices remain in the interval $[0,1]$; crucially, each full vertex in C is incident with precisely one upward and one downward open edge, so its weight remains 1. Moreover, we claim that f^*_λ is 0.9-balancing. In fact, for every $i \in [s]$ we have

$\mathrm{imb}(f^*_\lambda, i) = \mathrm{imb}(f^*, i)$. This can be seen by observing that every open edge in C that enters the clump $B_i \cup T_i$ either contributes an additional λ term to $f^*_\lambda(T_i)$ (if it is upward) or an additional $-\lambda$ term to $f^*_\lambda(B_i)$ (if it is downward) and so its added contribution to $f^*_\lambda(T_i) - f^*_\lambda(B_i)$ is λ. However, the next open edge along C leaves $B_i \cup T_i$ and, by similar reasoning, its added contribution to $f^*_\lambda(T_i) - f^*_\lambda(B_i)$ is $-\lambda$. The contributions cancel out. We conclude that $\mathrm{imb}(f^*_\lambda, i) = \mathrm{imb}(f^*, i)$, as claimed. As λ increases, eventually a point is reached where some open vertex or some open edge becomes closed. At that exact moment f^*_λ has fewer open vertices and/or open edges than f^*, contradicting the minimality of f^*. Therefore, H does not have cycles.

Since H is a non-empty forest, there exists a path P joining two distinct vertices of degree 1 in H, say x and y. Suppose that $x \in \bar{A}_i$, $y \in \bar{A}_j$. Since x and y have degree 1 in H, they are not full and they are not incident with fake edges, which means that x and y are the unique open vertices in their respective clumps $\bar{A}_i$ and $\bar{A}_j$. In particular, $i \neq j$ and P contains an open edge. Also, precisely one of $f^*(B_i)$ and $f^*(T_i)$ is an integer, and so $\mathrm{imb}(f^*, i) > 0$. Similarly, $\mathrm{imb}(f^*, j) > 0$. Like in the case where H had a cycle, we fix a direction for P and partition the open edges in P into *upward* and *downward* ones, depending on whether they go from $V^{(1)}$ to $V^{(2)}$ or the other way around. Let $\lambda \in \mathbb{R}$ be a number with small absolute value and define f^*_λ in the same way as previously, that is, by giving the upward edges of P additional weight λ and downward edges $-\lambda$. With the same reasoning as before, f^*_λ is a valid fractional matching provided that λ is small. Moreover, for every $m \in [s] \setminus \{i, j\}$ we have $\mathrm{imb}(f^*_\lambda, m) = \mathrm{imb}(f^*, m)$, also by an identical argument. However, the added contributions to $\mathrm{imb}(f^*_\lambda, i)$ and $\mathrm{imb}(f^*_\lambda, j)$ are non-zero. In fact, having the additional $\pm\lambda$ term either decreases or further increases the imbalance of the clumps $\bar{A}_i, \bar{A}_j$ by exactly $|\lambda|$. More precisely, there exist constants $s_i, s_j \in \{-1, 1\}$ such that, for small $|\lambda|$, $\mathrm{imb}(f^*_\lambda, i) = \mathrm{imb}(f^*, i) + s_i\lambda$ and $\mathrm{imb}(f^*_\lambda, j) = \mathrm{imb}(f^*, j) + s_j\lambda$. Therefore, for small $|\lambda|$, $\mathrm{imb}(f^*_\lambda) = \mathrm{imb}(f^*) + (s_i + s_j)\lambda$. Depending on the sign of $s_i + s_j$ we choose λ to be positive or negative, ensuring that $\mathrm{imb}(f^*_\lambda) \leq \mathrm{imb}(f^*) \leq 0.9$, which means that f^*_λ is 0.9-balancing. Finally, we keep increasing the magnitude of λ until some open vertex or open edge becomes closed. (Here it is important to note that the signs s_i, s_j cannot change before at least one open vertex become closed, at which time we stop our process.) This contradicts the minimality of f^*. Therefore, the auxiliary graph H is empty, and the claim follows. $\square$

Since all weights of f^* are 0 or 1, f^* gives rise to a matching M in G_σ. Furthermore, since f^* is 0-balancing, M balances $\bar{A}_1, \ldots, \bar{A}_s$. Lemma 16 follows. $\square$

5.3 Constructing the balancing paths in G

In this section we prove Lemma 5.

Proof of Lemma 5. The rough idea is as follows. We pull back a balancing matching M of $\bar{G}$, as given by Lemma 16, to G. The resulting subgraph $H_0 \subset G$ has maximum degree at most 2, is acyclic and 'overbalances' every β-almost-bipartite cluster A_i (the reason for this is that every β-almost-bipartite A_i gives rise to two clumps of $\bar{G}$, both of which are balanced by M; therefore, A_i gets balanced 'twice'). Since M is small, H_0 is also small, but it may have many components and, as a result, many leaves. To obtain property (c) in Lemma 5, in clusters with too many such vertices, we connect pairs of them by short paths. It turns out that in doing so we also fix the overbalancing issue. Therefore, we get properties (c) and (e) simultaneously. The remaining three properties are mainly technicalities. We use Proposition 6 to find the desired short paths in clusters.

Fix an ordering σ of $V(G)$ such that G_σ contains a matching M as given by Lemma 16; that is, M covers at most $|M| \le (\xi\zeta/4)n$ vertices and, for each $i \in [r]$, it satisfies $|T_i \setminus V(M)| = |B_i \setminus V(M)|$. Let H_0 be the subgraph of G spanned by edges $uv \in E(G)$ for which $u^{(1)}v^{(2)}$ or $v^{(1)}u^{(2)}$ is in M. By construction of G_σ, it is impossible for both $u^{(1)}v^{(2)}$ and $v^{(1)}u^{(2)}$ to be in M, and therefore $e(H_0) = e(M)$. Trivially, H_0 has no isolated vertices. Moreover, H_0 does not have cycles. Indeed, suppose to the contrary that H_0 contains a cycle $v_1 \dots v_\ell$. We may assume that $v_1^{(1)}v_2^{(2)}$ is in M. Since M is a matching, $v_2^{(2)}v_3^{(1)} \notin M$, and hence $v_2^{(1)}v_3^{(2)} \in M$. Similarly, $v_3^{(1)}v_4^{(2)}, \dots, v_{\ell-1}^{(1)}v_\ell^{(2)}, v_\ell^{(1)}v_1^{(2)}$ are edges in M. However, this implies that $\sigma(v_1) < \dots < \sigma(v_\ell) < \sigma(v_1)$, giving a contradiction.

We now show that the number of leaves of H_0 in A_i is even for every $i \in [r]$. For any subgraph $F \subset G$ and any set $U \subset V(G)$ we define $d_F(U) = \sum_{v \in U} d_F(v)$. We claim that $d_{H_0}(A_i)$ is even for every $i \in [r]$. Indeed, if A_i is γ-far-from-bipartite, then $d_M(A_i^{(1)}) = d_M(A_i^{(2)})$, as M balances the balanced bipartite graph with bipartition $\{A_i^{(1)}, A_i^{(2)}\}$, thus implying that $d_{H_0}(A_i) = d_M(A_i^{(1)}) + d_M(A_i^{(2)}) = 2d_M(A_i^{(1)})$. Now suppose that A_i is β-almost-bipartite and denote its prescribed bipartition by $\{X_i, Y_i\}$. Since M balances the two bipartite graphs with bipartitions $\{X_i^{(1)}, Y_i^{(2)}\}$ and $\{Y_i^{(1)}, X_i^{(2)}\}$, we have $d_{H_0}(X_i) - d_{H_0}(Y_i) = 2(|X_i| - |Y_i|)$. Either way, we see that $d_{H_0}(A_i)$ is even. Since all non-leaves in H_0 have degree 2, we find that the number of leaves of H_0 in A_i is even, as desired.

We proceed by extending H_0 to linear forests $H_0 \subset H_1 \subset \dots \subset H_m$ (for some $m \ge 0$) where, for each $j \in [m]$, H_j is obtained from H_{j-1} by adding a short path contained in some cluster A_i, joining two leaves of H_{j-1}. We stop when we reach a linear forest H_m that satisfies property (c).

Here is a more precise description of this process. Suppose that we have constructed linear forests $H_0 \subset H_1 \subset \dots \subset H_{j-1}$ where H_t contains an even number of leaves in A_i for every $t \in \{0, \dots, j-1\}$ and every $i \in [r]$. Suppose that H_{j-1} does not satisfy property (c). For convenience, we write $L = \{v \in V(H_{j-1}) : v \text{ is a leaf of } H_{j-1}\}$. We pick $i \in [r]$ such that $|A_i \cap L| \ne 0, 2$, so $|A_i \cap L| \ge 4$. Since every component of H_{j-1} is a path (and so contains two leaves), there exist vertices $x, y \in A_i \cap L$ that are in different components of H_{j-1}. By Proposition 6, $G[A_i]$ contains a path P_j of length at most $3/\zeta$, with ends x, y and whose vertex set does not intersect $V(H_{j-1}) \setminus \{x, y\}$. We set $H_j = H_{j-1} \cup P_j$ and note that our way of choosing x, y ensures that H_j is a linear forest. Moreover, since the set of leaves of H_j is the set of leaves of H_{j-1} minus $\{x, y\}$, the property that every cluster contains an even number of leaves still holds. This also implies that eventually we will find a linear forest H_m that satisfies property (c).

To justify the application of Proposition 6 in the previous paragraph, we note that, by our inductive construction, $|H_{j-1}| \le |H_0| + (3/\zeta)(j-1)$. Moreover, since H_{j-1} has $2(j-1)$ fewer leaves than H_0, we have $|H_0| - 2(j-1) \ge 0$, which implies that $j-1 \le |H_0|/2$, and therefore $|H_{j-1}| \le |H_0|(3/(2\zeta)+1) \le (\xi\zeta/4)(2/\zeta)n^2 \le (\zeta/6)n$, as needed.

It is clear that H_m satisfies properties (b) and (c). Also, by the same argument as above, $|H_m| \le (\xi\zeta/4)(2/\zeta)n \le (\xi/2)n$. We now focus on modifying H_m so that it also satisfies properties (d) and (e). Let $i \in [r]$ be the index of an arbitrary β-almost-bipartite cluster A_i and denote the prescribed bipartition of A_i by $\{X_i, Y_i\}$. First, suppose that X_i and Y_i have the same number t of leaves of H_m, so $t \in \{0,1\}$. Then $|V(H_m) \cap X_i| - |V(H_m) \cap Y_i| = (d_{H_m}(X_i)+t)/2 - (d_{H_m}(Y_i)+t)/2 = |X_i| - |Y_i|$, as $d_{H_0}(X_i) - d_{H_0}(Y_i) = 2(|X_i| - |Y_i|)$ (since M balances the two clumps corresponding to A_i) and $d_{H_j}(X_i) - d_{H_j}(Y_i)$ is the same for all $j \in [m]$, because for each $j \in [m]$, the path P_j that is added to H_{j-1} to form H_j is contained in one of the β-almost-bipartite clusters A_i, and thus $d_{P_j}(X_i) = d_{P_j}(Y_i)$ for each $i \in [r]$ such that A_i is β-almost-bipartite. It follows that properties (d) and (e) hold in this case. So, without loss of generality, we assume that both leaves of H_m are in X_i and denote them by x, x'. Since x has at least $(\delta/2)n > |H_m|$ neighbours in Y_i, it has a neighbour $y \in Y_i \setminus V(H_m)$. We define $e_i = xy$, with the intention of adding this edge to H_m to obtain the desired linear forest H. Clearly, $d_{H_m \cup \{e_i\}}(X_i) - d_{H_m \cup \{e_i\}}(Y_i) = d_{H_m}(X_i) - d_{H_m}(Y_i)$, x' is the unique leaf of $H_m \cup \{e_i\}$ in X_i and y is the unique such vertex in Y_i. The same calculation as in the previous case gives $|V(H_m \cup \{e_i\}) \cap X_i| - |V(H_m \cup \{e_i\}) \cap Y_i| = |X_i| - |Y_i|$.

The final definition of H is as follows: it is the subgraph of G spanned by the edges $E(H_m) \cup \{e_i : i \in [r] \text{ is such that } e_i \text{ is defined}\}$. It follows from the construction of H_m and the e_i's that H is a linear forest satisfying properties (b)-(e). Furthermore, $|H| \le |H_m| + r \le (\xi/2)n + r \le \xi n$, and so property (a) also holds. This completes the proof of Lemma 5. □

6 Concluding remarks

In this paper we prove that the vertices of every d-regular n-vertex graph, where $d \ge cn$ and $n \ge n_0(c)$, can be partitioned into at most $\lfloor n/(d+1) \rfloor$ cycles. It is natural to wonder whether this lower bound on d can be lowered. We believe that, with our methods, one could prove this result for $d \ge cn/\sqrt{\log\log n}$; we omit the details as we believe that the slight gain on the bound on d is not worth the loss in clarity of the proof. It would be interesting to obtain the result for smaller d, say, for $d \ge n^{1-\varepsilon}$ for some fixed $\varepsilon > 0$. We note that the main obstruction that prevents us from obtaining such a result is the first step of the proof, where the vertices are partitioned into well-behaved clusters. The reason for this is that we consider a sequence of parameters $\mu_1, \dots, \mu_r$, where μ_1 needs to be at most $\mu_r^{C^{r^2}}$ for some (large) fixed C. Since we also have that μ_r is bounded away from 1 and $\mu_1 \geq 1/n^2$, we find that $r = O(\sqrt{\log\log n})$, which implies that for our proof to work we must have $d = \Omega(n/\sqrt{\log\log n})$. On the other hand, the proof of Lemma 4 could be made to work for $d \geq n^{1-\varepsilon}$, similarly as in [25] (see also the remark at the end of Section 4). It would therefore be of interest to find a more efficient way to obtain the partition into well-behaved clusters.

It would also be interesting to determine if a version of our results holds for regular directed graphs or for regular oriented graphs.

Another possible direction for future research is to consider bipartite versions of the Bollobás and Häggkvist conjecture. Häggkvist [9] conjectured that every d-regular 2-connected bipartite graph on n vertices, where $d \geq n/6$ is Hamiltonian. This was essentially verified by Jackson and Li [13] who proved this statement for $d \geq (n+38)/6$. Recently, Li [24] conjectured that every d-regular 3-connected bipartite graph on n vertices, with $d \geq n/8$, is Hamiltonian. We suspect that our methods could be useful for this problem.

References

[1] J. Balogh, F. Mousset, and J. Skokan, *An extension of Dirac's Theorem*, Electron. J. Combin. **24** (2017), Paper #P3.56.

[2] B. Bollobás, *Extremal graph theory*, Academic Press, 1978, p167.

[3] B. Bollobás and A.M. Hobbs, *Hamiltonian cycles in regular graphs*, Advances in Graph Theory **3** (1978), 43–48.

[4] H. J. Broersma, J. van den Heuvel, B. Jackson, and H. J. Veldman, *Hamiltonicity of regular 2–connected graphs*, J. Graph Theory **22**, 105–124.

[5] L. DeBiasio and L. Nelsen, *Monochromatic cycle partitions of graphs with large minimum degree*, J. Combin. Theory Ser. B **122** (2017), 634–667.

[6] H. Enomoto, A. Kaneko, and Zs. Tuza, *Graphs of order n and minimum degree $n/4$*, Combinatorics (Proc. Coll. Math. Soc. János Bolyai, Eger, 1987), North-Holland, 1988, pp. 213–220.

[7] P. Erdős and A. M. Hobbs, *A class of Hamiltonian regular graphs*, J. Graph Theory **2** (1978), 129–135.

[8] G. Fan, *Longest cycles in regular graphs*, J. Combin. Theory Ser. B **39** (1985), 325–345.

[9] R. Häggkvist, *Unsolved problems*, Proc. Fifth Hungarian Colloq. Comb. (1976).

[10] J. Han, *On vertex-disjoint paths in regular graphs*, Electron. J. Combin. **25** (2018), Paper #P2.12.

[11] F. Hilbig, *Kantenstrukturen in nichthamiltonischen Graphen*, Ph.D. thesis, Technical University Berlin, 1986.

[12] B. Jackson, *Hamilton cycles in regular 2-connected graphs*, J. Combin. Theory Ser. B **29** (1980), 27–46.

[13] B. Jackson and H. Li, *Hamilton cycles in 2-connected regular bipartite graphs*, J. Combin. Theory B **62** (1994), 236–258.

[14] B. Jackson, H. Li, and Y. Zhu, *Dominating cycles in regular 3-connected graphs*, Discr. Math. **102** (1992), 163–176.

[15] H. A. Jung, *Longest circuits in 3-connected graphs*, Finite and infinite sets, Vol I, II, Colloq. Math. Soc. János Bolyai **37** (1984), 403–438.

[16] M. Kouider, *Covering vertices by cycles*, J. Graph Theory **18** (1994), 757–776.

[17] M. Kouider and Z. Lonc, *Covering cycles and k-term degree sums*, Combinatorica **16** (1996), 407–412.

[18] D. Kühn, A. Lo, D. Osthus, and K. Staden, *The robust component structure of dense regular graphs and applications*, Proc. London Math. Soc. **110**, 19–56.

[19] ———, *Solution to a problem of Bollobás and Häggkvist on Hamilton cycles in regular graphs*, J. Combin. Theory Ser. B **121** (2016), 85–145.

[20] D. Kühn and D. Osthus, *Hamilton partitions of regular expanders: a proof of Kelly's conjecture for large tournaments*, Adv. Math. **237** (2013), 62—146.

[21] D. Kühn, D. Osthus, and A. Treglown, *Hamiltonian degree sequences in digraphs*, J. Combin. Theory Ser. B **100** (2010), 367—-380.

[22] R. Kühn, D. Mycroft and D. Osthus, *A proof of Sumner's universal tournament conjecture for large tournaments*, Proc. London Math. Soc. **102** (2011), 731—-766.

[23] ———, *An approximate version of Sumner's universal tournament conjecture*, J. Combin. Theory Ser. B **101** (2011), 415—447.

[24] H. Li, *Generalizations of Dirac's theorem in Hamiltonian graph theory – a survey*, Discr. Math. **313** (2013), 2034–2053.

[25] A. Lo and V. Patel, *Hamilton cycles in sparse robustly expanding digraphs*, Electron. J. Combin. **25** (2018), no. 3, Paper #3.44.

[26] C. Magnant and D. Martin, *A note on the path cover number of regular graphs*, Australas. J. Combin. **43** (2009), 211—217.

[27] C. St. J. A. Nash-Williams, *Hamiltonian arcs and circuits*, Recent Trends in Graph Theory, Springer Berlin Heidelberg, 1971, pp. 197–210.

[28] Y. Zhu and H. Li, *Hamilton cycles in regular* 3*-connected graphs*, Discr. Math. **110** (1992), 229–249.

A Proof of Proposition 6

Proof of Proposition 6. Fix $R \subset V(H)$ with $|R| \leq (\zeta/6)|H|$ and let $x, y \in V(H) \setminus R$ be distinct vertices. We first observe that $H \setminus R$ is connected. Indeed, if $V(H) \setminus R$ admits a partition into non-empty sets X, Y with no $X-Y$ edges, then the number of $X-(Y \cup R)$ edges is at most $|X||R|$. We may assume that $|Y| \geq |X|$, and hence that $|Y \cup R| \geq |H|/2$, which implies that $|R| \leq (\zeta/3)|Y \cup R|$. However, this contradicts the assumption that the number of $X-(Y \cup R)$ edges in H is at least $\zeta|X||Y \cup R|$.

Now, we partition the vertices of $H \setminus R$ into sets according to their distance to x. That is, for all $i \geq 0$ we set

$$L_i = \{v \in V(H) \setminus R : \text{the shortest path from } x \text{ to } v \text{ in } H \setminus R \text{ has } i \text{ edges}\}$$

Since $H \setminus R$ is finite and connected, there exists a maximum value a for which L_a is non-empty and, for that value, $L_0, \ldots, L_a$ partition $V(H) \setminus R$.

Our aim is to show that $a \leq 3/\zeta$, so suppose that this is not the case. In particular, we have $a \geq 3$. Let j be an index in the set $[a-1]$ for which $|L_j|$ is minimal. We partition $V(H) \setminus R$ into two sets X, Y, defined by

$$\begin{cases} \text{if } j \geq \frac{a}{2}, \text{ then } X = L_0 \cup \cdots \cup L_j \text{ and } Y = L_{j+1} \cup \cdots \cup L_a, \\ \text{if } j < \frac{a}{2}, \text{ then } X = L_j \cup \cdots \cup L_a \text{ and } Y = L_0 \cup \cdots \cup L_{j-1}. \end{cases}$$

In either case X, Y are non-empty sets such that there are no edges between $X \setminus L_j$ and Y. Moreover, X contains at least $a/2$ of the sets $L_1, \ldots, L_{a-1}$, and so $|X| \geq |L_j|a/2$. Therefore, the number of $X-Y$ edges is at most $|L_j||Y| \leq (2/a)|X||Y|$.

We attach R to the larger one of the sets X, Y. For the following calculation we may assume that $|X| \geq |Y|$, in which case we consider the partition of $V(H)$ into sets $X \cup R, Y$. Since $|X \cup R| \geq |H|/2$ and $|R| \leq (\zeta/6)|H| \leq (\zeta/3)|X \cup R|$, the number of $R-Y$ edges is at most $(\zeta/3)|X \cup R||Y|$. Hence, the number of $(X \cup R)-Y$ edges does not exceed $(2/a)|X||Y| + (\zeta/3)|X \cup R||Y| \leq (2/a + \zeta/3)|X \cup R||Y|$. Therefore, we have $2/a + \zeta/3 \geq \zeta$, which implies that $a \leq 3/\zeta$, as desired. $\square$