Cycle-complete Ramsey numbers

Peter Keevash* Eoin Long† Jozef Skokan‡

July 16, 2018

Abstract

The Ramsey number $r(C_\ell, K_n)$ is the smallest natural number N such that every red/blue edge-colouring of a clique of order N contains a red cycle of length ℓ or a blue clique of order n. In 1978, Erdős, Faudree, Rousseau and Schelp conjectured that $r(C_\ell, K_n) = (\ell-1)(n-1)+1$ for $\ell \geq n \geq 3$ provided $(\ell, n) \neq (3,3)$.

We prove that, for some absolute constant $C \geq 1$, we have $r(C_\ell, K_n) = (\ell-1)(n-1)+1$ provided $\ell \geq C\frac{\log n}{\log\log n}$. Up to the value of C this is tight since we also show that, for any $\varepsilon > 0$ and $n > n_0(\varepsilon)$, we have $r(C_\ell, K_n) \gg (\ell-1)(n-1)+1$ for all $3 \leq \ell \leq (1-\varepsilon)\frac{\log n}{\log\log n}$.

This proves the conjecture of Erdős, Faudree, Rousseau and Schelp for large ℓ, a stronger form of the conjecture due to Nikiforov, and answers (up to multiplicative constants) two further questions of Erdős, Faudree, Rousseau and Schelp.

1 Introduction

Graph Ramsey numbers are a central topic of research in Combinatorics. Given two graphs G and H, the Ramsey number $r(G,H)$ is the smallest natural number N such that every red/blue colouring of the edges of the complete graph K_N on N vertices contains a red copy of G or a blue copy of H. The existence of $r(G,H)$ follows from Ramsey's theorem [42], but determining or accurately estimating these parameters presents many challenging problems.

The classical Ramsey numbers are the graph Ramsey numbers $r(G,H)$ where G and H are cliques. Erdős and Szekeres [23] showed $r(K_n,K_n) \leq 2^{(1+o(1))2n}$, and later Erdős [20] showed $r(K_n,K_n) \geq 2^{(1+o(1))n/2}$, in one of the first instances of the probabilistic method. Both bounds changed very little over the past 70 years, despite progress by Thomason [51] and Conlon [17] on the upper bound, and by Spencer [48] on the lower bound. Another intensively studied Ramsey number is $r(K_3,K_n)$; it was a long-standing open problem to determine its order of magnitude, which is now known to be $\Theta\big(\frac{n^2}{\log n}\big)$, due to theorems of Ajtai, Komlós and Szemerédi [4] and Kim [31]. Recent analyses of the triangle-free process independently by Bohman and Keevash [5] and by Fiz Pontiveros, Griffiths and Morris [25], together with an improved upper bound due to Shearer [46], have now determined $r(K_3,K_n)$ to within a multiplicative factor of $4+o(1)$.

*Mathematical Institute, University of Oxford, Oxford, UK. E-mail: keevash@maths.ox.ac.uk. Research supported in part by ERC Consolidator Grant 647678.

†Mathematical Institute, University of Oxford, Oxford, UK. E-mail: long@maths.ox.ac.uk.

‡Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK, *and* Department of Mathematics, University of Illinois, 1409 W. Green Street, Urbana IL 61801, USA. E-mail: j.skokan@lse.ac.uk. Research supported in part by NSF Grant DMS-1500121.

At the other end of the spectrum, sparse graphs tend to have small Ramsey numbers. In this context, Chvátal, Rödl, Szemerédi and Trotter [16] proved that if G and H have bounded maximum degree then $r(G,H) = O(v(G) + v(H))$, where $v(G)$ denotes the number of vertices of the graph G. A similar bound was obtained by Chen and Schelp [13] under the assumption of bounded arrangeability. After intense effort [2, 26, 27, 33, 34, 35], a longstanding conjecture of Burr and Erdős [10] that such bounds hold only assuming bounded degeneracy was recently confirmed by Lee [36].

In this paper, we will focus on the cycle-complete Ramsey numbers $r(C_\ell, K_n)$. For any connected graph H, Chvátal and Harary [15] observed that $r(H,K_n) \geq (v(H)-1)(n-1)+1$. This is shown by the red/blue edge-coloured clique of order $(v(H)-1)(n-1)$, in which the red edges consist of $n-1$ disjoint cliques of order $v(H)-1$ and all the remaining edges are blue. Burr and Erdős [11] asked when equality holds in the Chvátal–Harary bound (the 'Ramsey goodness' question, see e.g. [1]). When $H = C_\ell$, for $\ell \geq n^2 - 2$ Bondy and Erdős [6] showed the equality

$$r(C_\ell, K_n) = (\ell-1)(n-1)+1. \tag{1}$$

Erdős, Faudree, Rousseau and Schelp [21] noted that $r(C_3, K_n) = r(K_3, K_n)$ grows much faster than a linear function of n (as discussed above), and posed the problem of determining the critical ℓ at which the change in behaviour of $r(C_\ell, K_n)$ occurs. They conjectured (see also [14, Chapter 2]) that (1) holds for $\ell \geq n \geq 3$ provided $(\ell, n) \neq (3,3)$.

There is a large literature on $r(C_\ell, K_n)$. An improved lower bound on $r(C_\ell, K_n)$ for small ℓ was given by Spencer [47]. Caro, Li, Rousseau and Zhang [12] improved the upper bound on $r(C_\ell, K_n)$ of Erdős et al. [21] for small even ℓ; Sudakov [49] gave a similar improvement for small odd ℓ. Several authors [24, 43, 52, 8, 44] confirmed the Erdős–Faudree–Rousseau–Schelp conjecture for small values of n. Schiermeyer [45] improved the result of Bondy and Erdős by showing that (1) holds for $\ell \geq n^2 - 2n > 3$. Nikiforov [40] substantially extended this range, proving that (1) holds for $\ell \geq 4n+2$. Moreover, he conjectured (Conjecture 2.14 in [40]) that in fact (1) already holds at a much lower threshold, namely that for all $\varepsilon > 0$ there is n_0 such that $r(C_\ell, K_n) = (\ell-1)(n-1)+1$ provided $\ell \geq n^\varepsilon$ and $n \geq n_0$.

Our main result proves both the Erdős–Faudree–Rousseau–Schelp conjecture for large ℓ and Nikiforov's conjecture. In fact, we prove (1) for a much wider range of parameters.

Theorem 1.1. *There is $C \geq 1$ so that $r(C_\ell, K_n) = (\ell-1)(n-1)+1$ for $n \geq 3$ and $\ell \geq C\frac{\log n}{\log\log n}$.*

Remarks: All logarithms in this paper are to base 2. Note that $r(C_\ell, K_1) = 1$ and $r(C_\ell, K_2) = \ell$ for all $\ell \geq 3$; we include the condition $n \geq 3$ only to avoid division by 0 in the lower bound on ℓ.

The bound in Theorem 1.1 is best possible up to the value of C, as shown by our next result.

Theorem 1.2. *Given $\varepsilon > 0$ there is $n_0(\varepsilon)$ so that $r(C_\ell, K_n) > n\log n \gg (\ell-1)(n-1)+1$ for all $n \geq n_0(\varepsilon)$ and $3 \leq \ell \leq (1-\varepsilon)\frac{\log n}{\log\log n}$.*

In combination, Theorems 1.1 and 1.2 answer (up to the constant C) two further questions of Erdős et al. [21] regarding $r(C_\ell, K_n)$, namely (i) the location of the critical value of ℓ for the transition in behaviour of $r(C_\ell, K_n)$, and (ii) the choice of ℓ that minimises $r(C_\ell, K_n)$. The answer to both questions is $\ell = \Theta\big(\frac{\log n}{\log\log n}\big)$.

An overview of the proof of Theorem 1.1 and the organisation of the paper is as follows. We suppose for a contradiction that there is some C_ℓ-free graph G with $v(G) = N = (\ell-1)(n-1)+1$ and independence number $\alpha(G) \leq n-1$. By induction we can also assume G has minimum degree $\delta(G) \geq \ell - 1$. The main task of the paper is to prove the stability result (Lemma 5.1) that G is close in structure to the lower bound construction described above, i.e. G can be mostly partitioned into approximate cliques of size about ℓ (and also less than ℓ, as there is no C_ℓ). Then in Section 6, following various arguments to clean up the approximate structure, we will see that it is incompatible with our assumptions, and so obtain a contradiction that proves the theorem.

In the next section, after the short proof of Theorem 1.2, we gather various tools needed for the proof of the stability result. Over the following three sections we prove the existence of approximate decompositions of G into pieces whose properties are gradually strengthened: in Section 3 the pieces are quite dense, in Section 4 they are 'hubs' (highly connected in a certain sense), and in Section 5 they are 'almost cliques', as required for the stability result.

2 Preliminaries

We start in the next subsection with some notation, then we prove Theorem 1.2. In the third subsection we collect various well-known results that we use in our proofs. The final subsection of this section describes two applications of Breadth First Search.

2.1 Notation

We summarise some (mostly) standard graph theory notation (see e.g. [7]) used in this paper. Let G be a finite graph. We write $v(G) := |V(G)|$ for the number of vertices and $e(G) := |E(G)|$ for the number of edges. Given a vertex $v \in V(G)$, the neighbourhood of v in G is $N_G(v) := \{y \in V(G) : xy \in E(G)\}$. The degree of v is $d_G(v) := |N_G(v)|$. The minimum degree is $\delta(G) := \min\{d(v) : v \in V(G)\}$, the maximum degree is $\Delta(G) := \max\{d(v) : v \in V(G)\}$, and the average degree is $d(G) := 2e(G)/v(G)$. Given $A \subset V(G)$, the induced graph $G[A]$ has vertex set A and edge set $\{e \in E(G) : e \subset A\}$. Given disjoint sets $A, B \subset V(G)$, we let $G[A, B]$ denote the bipartite graph with parts A and B and edge set $\{e \in E(G) : |e \cap A| = |e \cap B| = 1\}$. A path $P = x_0x_1 \dots x_\ell$ of length ℓ consists of $\ell + 1$ distinct vertices $x_0, \dots, x_\ell$, where x_ix_{i+1} is an edge for $i \in \{0, \dots, \ell - 1\}$. We call x_0 and x_k the end vertices of P and say that P is an x_0x_k-path. We say P is internally disjoint from a set X if X contains none of the interior vertices $\{x_1, \dots, x_{\ell-1}\}$ of P. A cycle of length ℓ, or ℓ-cycle, is a graph obtained from a path $P = x_0x_1 \dots x_{\ell-1}$ of length $\ell - 1$ by adding the edge $x_{\ell-1}x_0$. Edges of cycles will often be listed modulo ℓ, so that $x_{\ell-1}x_\ell$ represents $x_{\ell-1}x_0$. We say $I \subset V(G)$ is independent if $G[I]$ has no edges. The independence number $\alpha(G)$ is the size of a largest independent set in G. Given natural numbers $m \le n$ we let $[m, n] := \{m, m+1, \dots, n\}$. To simplify the presentation, we may omit floor and ceiling signs when they are not crucial.

2.2 The lower bound

The lower bound construction comes from the following application of the probabilistic method.

Proof of Theorem 1.2. Let $\varepsilon \in (0, 1)$, $n > n_0(\varepsilon)$ and $N = 2n \log n$. It suffices to prove that there is a graph G on at least $N/2$ vertices with $\alpha(G) < n$ which does not contain a cycle C_ℓ with $\ell \le \ell_0 := (1 - \varepsilon)\frac{\log n}{\log\log n}$. We consider a random graph $G_1 \sim G(N, p)$, where $p := \frac{3\log\log n}{n-1}$. The expected number of independent sets of order n in G_1 is

$$\binom{N}{n}(1-p)^{\binom{n}{2}} \le \Big(\frac{eN}{n}e^{-p(n-1)/2}\Big)^n = \big(2e(\log n)e^{-3\log\log n/2}\big)^n \ll 1.$$

On the other hand, the expected number of cycles of length at most ℓ_0 is $\sum_{i\in[3,\ell_0]}(Np)^i \le 2(Np)^{\ell_0} \le 2(7\log n \log\log n)^{(1-\varepsilon)\log n/\log\log n} \ll N$. By Markov's inequality applied to both of these expectations, with positive probability G_1 satisfies $\alpha(G_1) < n$ and has $\leq N/2$ cycles of length at most ℓ_0. Fixing a choice of such G_1 and deleting a vertex from each cycle of length at most ℓ_0 leaves a graph G with the required properties. $\square$

2.3 Tools

In this subsection we collect several well-known results. The first is very simple, but we include a short proof for the convenience of the reader.

Proposition 2.1. *For any graph G,*

(i) if G has no subgraph of minimum degree at least k then $e(G) \leq \binom{k}{2} + (v(G) - k)(k-1)$;

(ii) G contains a subgraph G_1 with $\delta(G_1) \geq d(G)/2$;

(iii) G contains a bipartite subgraph G_2 with $d(G_2) \geq d(G)/2$.

Proof. To see (i), note that as any subgraph of G contains a vertex with degree at most $k-1$, we may iteratively delete such vertices until we obtain a subgraph on k vertices. The bound follows by counting edges. Similarly, for (ii), if there were no such G_1 we could reduce G to an empty graph by deleting vertices of degree less than $d(G)/2$, but then $e(G) < v(G)d(G)/2$ would be a contradiction. Lastly, for (iii), note that a random induced bipartite subgraph G_2 of G has $\mathbb{E}d(G_2) = d(G)/2$. $\square$

Next we state several classical results from extremal graph theory.

Theorem 2.2 (Turán [50])**.** *Any graph G satisfies $\alpha(G) \geq \frac{v(G)}{d(G)+1}$.*

Theorem 2.3 (Dirac [18])**.** *Any graph G with $d(G) \geq v(G)/2$ contains a Hamilton cycle.*

Theorem 2.4 (Bondy [9])**.** *Any graph G with $d(G) \geq v(G)/2$ is either a complete bipartite graph or is pancyclic, i.e. contains cycles of all lengths in $[3, v(G)]$.*

Theorem 2.5 (Erdős and Gallai [22])**.** *Any graph G with $d(G) > k-1$ has a path of length k.*

We conclude by stating a version of Dependent Random Choice (see [28, Lemma 7.2]).

Theorem 2.6. *Given $\varepsilon > 0$ there is $\delta > 0$ so that the following holds for $N \geq N_0(\varepsilon)$ and any N-vertex graph G with at least $N^{2-\delta}$ edges. There are disjoint sets $U_1, U_2 \subset V(G)$ such that, for $i = 1, 2$, every $a, a' \in U_i$ satisfies $|N_G(a, U_{3-i}) \cap N_G(a', U_{3-i})| \geq N^{1-\varepsilon}$.*

2.4 Breadth First Search

Here give two applications of Breadth First Search, namely finding short cycles, and a nice decomposition of a substantial part of any graph.

We start by describing the well-known construction of a breadth first search tree T in a graph G rooted at some vertex $x \in V(G)$. At each step $i \geq 0$, we construct a tree T_i with layers $V_0, \ldots, V_i$ which are disjoint subsets of $V(G)$. Initially, T_0 is a tree with one vertex, namely $V(T_0) = V_0 = \{x\}$. Given T_{i-1} for some $i > 0$, we let $V_i := N_G(V_{i-1}) \setminus V(T_{i-1})$. If $V_i = \emptyset$ we terminate with $T = T_{i-1}$, otherwise we obtain T_i from T_{i-1} by adding an arbitrary edge of G from each vertex in V_i to some vertex in V_{i-1}. It will be useful to consider the first layer which does not cause the tree to grow significantly, in the sense of the following simple proposition.

Proposition 2.7. *Let $\gamma > 1$ and let G be an N-vertex graph. Let T be a breadth first search tree in G rooted at $x \in V(G)$ with layers $V_0, \ldots, V_r$. Suppose $m \in \mathbb{N}$ is minimal such that $|\cup_{i=0}^{m+1} V_i| \leq \gamma |\cup_{i=0}^{m} V_i|$. Then $m \leq \frac{\log N}{\log \gamma} = \log_\gamma(N)$.*

Proof. By definition of m we have $N \geq |\bigcup_{i\in[m]} V_i| \geq \gamma|\bigcup_{i\in[m-1]} V_i| \geq \ldots \geq \gamma^m|V_0| = \gamma^m$. $\square$

Our first application is to finding short cycles within an approximate range.

Lemma 2.8. *Let G be an N-vertex graph with $d(G) \geq d = 16\gamma d_1$, where $\gamma > 1$ and $d_1 \geq 2$. Then G contains an ℓ-cycle for some $\ell \in \big[d_1, d_1 + 2\log_\gamma(N)\big]$.*

Proof. By Lemma 2.1 (ii) and (iii) there is a bipartite subgraph G' of G with $\delta(G') \geq d(G)/4$. Let T be a breadth first search tree in G' rooted at some $x \in V(G')$ with layers $V_0, \ldots, V_r$. Let $m \leq \log_\gamma(N)$ be as in Proposition 2.7. As G' is bipartite, we have $G'[V_i] = \emptyset$ for all $i \in [r]$, so

$$\sum_{i\in[0,m]} e(G'[V_i, V_{i+1}]) = e\big(G'[\cup_{i\in[0,m+1]}V_i]\big) \geq \frac{\delta(G')}{2}\sum_{i\in[0,m]}|V_i| \geq \sum_{i\in[0,m]}\frac{d(G)}{16\gamma}\big(|V_i| + |V_{i+1}|\big),$$

using $\sum_{i\in[0,m]}(|V_i| + |V_{i+1}|) \leq (1+\gamma)\sum_{i\in[0,m]}|V_i| \leq 2\gamma\sum_{i\in[0,m]}|V_i|$. Thus $d(G'[V_i, V_{i+1}]) \geq d(G)/16\gamma \geq d_1$ for some $i \in [m]$. By Theorem 2.5, there is a path of length d_1 in $G[V_i, V_{i+1}]$. By possibly removing vertices we can obtain an xy-path in $G[V_i, V_{i+1}]$ of length between $d_1 - 2$ and d_1 with $x, y \in V_i$. Combining this with the unique xy-path in T of length at most $2m \leq 2\log_\gamma(N)$ gives a cycle of length in $\big[d_1, d_1 + 2\log_\gamma(N)\big]$, as required. $\square$

Our second application is to construct the following partial decomposition of a graph G, consisting of a family of disjoint sets $X_i \subset V(G)$, which are mutually non-adjacent in G, with each X_i entirely at a fixed distance in some tree T_i from the root x_i.

Lemma 2.9. *Let $\gamma > 1$ and let G be an N-vertex graph. Then there are triples $\{(x_i, X_i, T_i)\}_{i\in[t]}$, where each T_i is a subtree of G rooted at x_i and $X_i \subset V(T_i)$, such that:*

(i) there is $d_i \in [0, \log_\gamma(N)]$ such that for all $x'_i \in X_i$ the unique $x_i x'_i$-path in T_i has length d_i;

(ii) $\{X_i\}_{i\in[t]}$ are disjoint and satisfy $|\bigcup_{i\in[t]} X_i| \geq N/2\gamma$;

(iii) there are no edges of G between X_i and X_j for distinct $i, j \in [t]$.

Proof. We prove the statement by induction on $v(G)$, noting that it is trivial if $v(G) = 1$.

Let T be a breadth first search tree in G rooted at some $x \in V(G)$ with layers $V_0, \ldots, V_r$. Let $m \leq \log_\gamma(N)$ be as in Proposition 2.7. Let $\{X_i\}_{i\in[s]}$ be $\{V_{2i}\}_{2i\in[m]}$ or $\{V_{2i+1}\}_{2i+1\in[m]}$ according to which set $\bigcup_{2i\in[m]} V_{2i}$ or $\bigcup_{2i+1\in[m]} V_{2i+1}$ is larger. Setting $X := \bigcup_{i\in[s]} X_i$, we note that $|X| + |N_G(X)| = |X \cup N_G(X)| \leq |\bigcup_{i\in[m+1]} V_i| \leq \gamma|\bigcup_{i\in[m]} V_i| \leq 2\gamma|X|$.

For each $i \in [s]$, set $x_i = x$, $T_i = T$ and $d_i = j$, where $X_i = V_j$, so that (i) holds by the definition of V_j. As $\{X_i\}_{i\in[s]}$ are non-consecutive layers of a breadth first search tree, they are disjoint and there are no edges between X_i and X_j for distinct $i, j \in [s]$.

Now let $W = V(G)\setminus(X \cup N_G(X))$ and apply induction on $G[W]$ to obtain $\{(x_i, X_i, T_i)\}_{i\in[s+1,t]}$. We claim that $\{(x_i, X_i, T_i)\}_{i\in[t]}$ satisfy the statement of the lemma. Indeed, (i) holds by construction. For (ii), disjointness is clear, and we have $\sum_{i\in[s]}|X_i| = |X| \geq |X \cup N_G(X)|/2\gamma$ and $\sum_{i\in[s+1,t]}|X_i| \geq |W|/2\gamma$ by induction. Finally, (iii) holds by construction and as each X_j with $j \in [s+1, t]$ is contained in W, which is disjoint from $X \cup N_G(X)$. $\square$

3 Quite dense subgraphs

In this section we take our first steps towards the stability result described above, by showing that any supposed counterexample to Theorem 1.1 can be partitioned almost entirely into vertex-disjoint subgraphs, each of which is quite large (has $\ell^{1-o(1)}$ vertices) and is quite dense (has $\ell^{2-o(1)}$ edges). We start by showing that a graph of large minimum degree has a long path or a dense subgraph.

Lemma 3.1. *Fix $D \in \mathbb{N}$. Then any graph G either*

(i) contains paths of length at least D starting at any given vertex, or

(ii) has a subgraph H with $v(H) \le D$ and $e(H) \ge \binom{\delta(G)+1}{2}$.

Proof. Suppose that (i) fails, i.e. there is $x_0 \in V(G)$ such that any path starting at x_0 has length less than D. We must show that (ii) holds. We construct a path P starting at x_0 as follows. At step $i \ge 0$, having chosen a path $P_{i-1} = x_0 \dots x_{i-1}$, we select $x_i \in N_G(x_{i-1}) \setminus \{x_0, \dots, x_{i-1}\}$ that maximises $|N_G(x_i) \cap \{x_0, \dots, x_{i-1}\}|$. If no such x_i exists we terminate with $P = P_{i-1}$. Let $P = x_0 x_1 \cdots x_\ell$ be the final path, where by choice of x_0 we have $\ell < D$. By the termination rule, we have $N_G(x_\ell) \subset V(P)$. Let $N_G(x_\ell) = \{x_{i_1}, \dots, x_{i_s}\}$, where $s \ge \delta(G)$, ordered so that $i_1 < \dots < i_s$. As x_ℓ is adjacent to x_{i_j} for each $j \in [s]$, the rule for choosing x_{i_j} guarantees $|N_G(x_{i_j+1}) \cap \{x_0, \dots, x_{i_j}\}| \ge |N_G(x_\ell) \cap \{x_0, \dots, x_{i_j}\}| = j$ for each $j \in [s]$. Then $H = G[V(P)]$ satisfies $v(H) \le D$ and $e(H) \ge \sum_{j=1}^{s} j \ge \binom{\delta(G)+1}{2}$. □

Remark: An unpublished result of the second author in [38] used a variant of Lemma 3.1 to prove that subgraphs of the cube graph with average degree d contain paths and cycles of length at least $2^{\Omega(\sqrt{d})}$. This result was later improved to $2^{\Omega(d)}$ in [37] via a different approach.

We combine the previous lemma with two applications of the breadth first search decomposition of the previous section to show that any C_ℓ-free graph with small independence number contains a small dense subgraph.

Lemma 3.2. *Let $N, D, \ell \in \mathbb{N}$, $\gamma > 1$, where $3\log_\gamma(N) \le \ell \le D$, and $d \ge 8\gamma^2$. Suppose G is a C_ℓ-free graph on N vertices with $\alpha(G) \le N/d$. Then G has a subgraph H with $v(H) \le D$ and $e(H) \ge d^2/2^9\gamma^4$.*

Proof. Let $\{(x_i, X_i, T_i)\}_{i\in[t]}$ be obtained by applying Lemma 2.9 to G. Let $X = \bigcup_{i\in[t]} X_i$, and note that $|X| \ge N/2\gamma$. Let $\{(y_i, Y_i, T'_i)\}_{i\in[s]}$ be obtained by applying Lemma 2.9 again, this time to $G[X]$. Let $Y = \bigcup_{i\in[s]} Y_i$, and note that $|Y| \ge |X|/2\gamma \ge N/4\gamma^2 \ge d\alpha(G)/4\gamma^2$. By Theorem 2.2 (Turán's Theorem), $d(G[Y]) \ge d/4\gamma^2 - 1 \ge d/8\gamma^2$, as $d \ge 8\gamma^2$. Then Proposition 2.1 (ii) applied to $G[Y]$ gives some $G' = G[Y']$ with $Y' \subset Y$ such that $\delta(G') \ge d/16\gamma^2$. By Lemma 3.1, to complete the proof of the lemma, it suffices to show that G' does not contain a path of length D.

For contradiction, suppose $P = z_0 z_1 \dots z_D$ is a path in G'. As $z_0 \in Y$ there is a triple (y_j, Y_j, T'_j) with $z_0 \in Y_j$. As T'_j is a tree, and so a connected subgraph of $G[X]$, by Lemma 2.9 (iii) there is a triple (x_i, X_i, T_i) with $V(T'_j) \subset X_i$, and by (i) there is $d_i \in [0, \log_\gamma(N)]$ so that every vertex in X_i is at distance d_i from x_i in T_i. In particular, the $x_i y_j$-path and $x_i z_0$-path in T_i only intersect X_i in y_j and z_0. We let P_1 be the $y_j z_0$-path in T_i. Then P_1 has length $\ell_1 \le 2\log_\gamma(N)$ and intersects X_i only in y_j and z_0.

We now use the triple (y_j, Y_j, T'_j). As P is a connected subgraph of $G[Y]$, by Lemma 2.9 (iii) we have $V(P) \subset Y_j$, and by (i) there is $d'_j \in [0, \log_\gamma(N)]$ so that every vertex of Y_j is at distance d'_j from y_j in T'_j. Let $\ell_2 = \ell - \ell_1 - d'_j$ and consider the subpath $P_2 = z_0 z_1 \dots z_{\ell_2}$ of P. Let P_3 be the $y_j z_{\ell_2}$-path in T'_j. Then P_3 has length d'_j and intersects Y_j only in z_{ℓ_2}. As $V(P_3) \subset V(T'_j) \subset X_i$, we can combine P_1, P_2 and P_3 to form a cycle of length ℓ. This contradiction completes the proof. □

By iterating the previous lemma one can obtain the following approximate decomposition of the vertex set of G. This Corollary will not be used in the proof of Theorem 1.1 so we omit its proof, which is similar to that of Corollary 4.3 in the next section.

Corollary 3.3. *Given $\varepsilon > 0$ there is $C \geq 1$ so that the following holds for all $\ell, n \in \mathbb{N}$ with $n \geq 3$ and $\ell \geq C\frac{\log n}{\log\log n}$. Suppose G is a C_ℓ-free graph on $N = (\ell-1)(n-1)+1$ vertices with $\alpha(G) \leq n-1$. Then there is a partition $V(G) = W \cup \bigcup_{i\in[L]} V_i$ so that $|V_i| < \ell$ and $e(G[V_i]) > \ell^{2-\varepsilon}$ for all $i \in [L]$, and $|W| \leq \varepsilon N$.*

4 Hubs

Continuing our progress towards the stability result, we next upgrade the properties of our decomposition by showing that the quite dense pieces from the last section must contain quite large 'hubs', which have the property that any small set of vertices can be joined together via disjoint paths of essentially any desired lengths. The precise definition is as follows.

Definition 4.1. *Let G be a graph and $A, B \subset V(G)$ be disjoint sets. Given distinct $x, y \in A \cup B$, we call $\ell \in \mathbb{N}$ a* bipartite length *for $\{x, y\}$ in $G[A, B]$ if (a) ℓ is even and $\{x, y\} \subset A$ or $\{x, y\} \subset B$, or (b) ℓ is odd and $|\{x, y\} \cap A| = |\{x, y\} \cap B| = 1$.*

For $\varepsilon \in (0, 1)$ and $u \in \mathbb{N}$, we call a triple (A, B, D) an (u, ε)-hub in a graph G if $|A| = |B| = u$, $|D| \leq \varepsilon u$, and for any distinct $s_1, \ldots, s_m, t_1, \ldots, t_m$ in $A \cup B$ with $m \leq u^{1-\varepsilon}$ we have the following connection property: for any $\ell_1, \ldots, \ell_m \geq 2$ with $\sum_{i\in[m]}(\ell_i + 1) \leq 2(1-\varepsilon)u$, where each ℓ_i is a bipartite length for $\{s_i, t_i\}$ in $G[A, B]$, there are vertex-disjoint paths $P_1, \ldots, P_m$ in $G[A \cup B \cup D]$, where each P_i is an $s_i t_i$-path of length ℓ_i.

The main lemma of this section shows that quite dense graphs contain large hubs.

Lemma 4.2. *Given $\varepsilon \in (0, 1)$ there is $\delta > 0$ so that for $N \geq N_0(\varepsilon)$ and any integer $u \in [N^\varepsilon, N^{1-\varepsilon}]$, every N-vertex graph G with $d(G) \geq N^{1-\delta}$ contains a (u, ε)-hub.*

Proof. We assume throughout the proof that δ is sufficiently small and N is sufficiently large. By Proposition 2.1 (iii) we may assume G is bipartite. Let $\delta_1 = \varepsilon^2/10$. By Theorem 2.6, applied with δ_1 in place of ε, there are disjoint $U_1, U_2 \subset V(G)$ so that $|N_G(a, U_{3-i}) \cap N_G(a', U_{3-i})| \geq N^{1-\delta_1}$ for every $a, a' \in U_i$ with $i \in [2]$. As G is bipartite, U_1 and U_2 must lie on opposite sides of the bipartition. We construct an alternating cycle C of length $2u$ in $G[U_1, U_2]$ by fixing distinct vertices $a_1, \ldots, a_u \in U_1$ and greedily selecting a common neighbour in U_2 of each consecutive pair $\{a_i, a_{i+1}\}$ (including $\{a_u, a_1\}$) so that all selected vertices are distinct. This is possible as $u \leq N^{1-\varepsilon} \ll N^{1-\delta_1}$. We let $A = V(C) \cap U_1 = \{a_1, \ldots, a_u\}$ and $B = V(C) \cap U_2$.

We let D be a random subset of $(U_1 \cup U_2) \setminus (A \cup B)$ where each element is included independently with probability $p = \varepsilon u/2N$. By Markov's inequality, $|D| \leq 2pN \leq \varepsilon u$ with probability at least $1/2$. Furthermore, for each pair $a, a' \in A$, we have

$$\mathbb{E}\big(|N_G(a) \cap N_G(a') \cap D|\big) \geq p\big(|N_G(a, U_2) \cap N_G(a', U_2)| - |U_2 \cap C|\big) \geq \varepsilon u/4N^{\delta_1} \geq 2u^{1-\varepsilon/2},$$

and similarly for each pair in B. By Chernoff's inequality (see [3, Appendix A]), with positive probability D satisfies $|D| \leq \varepsilon u$ and $|N_G(c) \cap N_G(c') \cap D| \geq u^{1-\varepsilon/2}$ for all $\{c, c'\} \subset A$ or $\{c, c'\} \subset B$. We fix any set D with these properties.

It remains to show that (A, B, D) is a (u, ε)-hub. Suppose $S = \{s_1, \ldots, s_m\}$ and $T = \{t_1, \ldots, t_m\}$ are disjoint subsets of $A \cup B$ with $m \leq u^{1-\varepsilon}$. Let $\ell_1, \ldots, \ell_m \geq 2$ with $\sum_{i\in[m]}(\ell_i + 1) \leq 2(1-\varepsilon)u$, where each ℓ_i is a bipartite length for $\{s_i, t_i\}$ in $G[A, B]$. We want to find vertex-disjoint paths $P_1, \ldots, P_m$ in $G[A \cup B \cup D]$, where each P_i is an $s_i t_i$-path of length ℓ_i.

First we claim that there is a path R with $V(R) \cap (S \cup T) = \emptyset$, $|V(R) \cap D| \leq 2m$ and $|(A \cup B) \setminus V(R)| \leq 4m$. To see this, we consider $C \setminus (S \cup T)$, which is the vertex-disjoint union of some paths $R_1, \dots, R_k$, where $k \le 2m$. By deleting at most two vertices from each such path R_i, we can assume that each starts and ends in A. We form R by 'stitching' these paths together greedily, using distinct vertices from $D \cap B$ to link successive paths R_i and R_{i+1} for all $i \in [k-1]$. This is possible by the common neigbourhood property, as $2m \ll u^{1-\varepsilon/2}$, so the claim follows.

Now we will construct the paths $P_1, \dots, P_m$ by chopping R into suitable subpaths and connecting these to the endpoint sets S and T. To construct P_1, we consider separately the cases $\ell_1 = 2$, $\ell_1 = 3$ and $\ell_1 \ge 4$. If $\ell_1 = 2$ we let $P_1 = s_1u_1t_1$ for any common neighbour $u_1 \in D$ of s_1 and t_1 disjoint from all previous choices. If $\ell_1 = 3$ we let $P_1 = s_1u_1v_1t_1$ where $u_1 \in D$ is a neighbour of s_1 and $v_1 \in D$ is a common neighbour of t_1 and u_1, with $\{u_1, v_1\}$ disjoint from all previous choices. Lastly, if $\ell_1 \ge 4$ we consider a subpath R_1 starting at one end of R with length $\ell_1 - 3$. As ℓ_1 is a bipartite length for $\{s_1, t_1\}$, it is possible to delete a vertex from one end of R_1 to obtain a subpath R_1' of length $\ell_1 - 4$ which starts on the same side of the partition as s_1 and ends on the same side as t_1. Writing x_1 and y_1 for the ends of R_1', we form the s_1t_1-path P_1 of length ℓ_1 from R_1' by adding paths $s_1u_1x_1$ and $t_1v_1y_1$ where $u_1 \in D$ is a common neighbour of s_1 and x_1, and $v_1 \in D$ is a common neighbour of t_1 and y_1, with $\{u_1, v_1\}$ disjoint from all previous choices. To continue, we modify R by removing R_1, then repeat the process to find P_2, and so on.

It remains to show that the above process succeeds, i.e. that we do not ever exhaust R or any common neighbourhoods in D. To see this, note that initially $|R| \ge |\bigcup_{i\in[k]} R_i| \ge 2u - |S \cup T| - 2m \ge \sum_{i\in[m]} \ell_i$. As we remove at most ℓ_i vertices from R to build each path P_i, we never run out of vertices in R. Also, we used at most 2 vertices from D to build each P_i, and so at most $4m \le u^{1-\varepsilon/2}/2$ from D in total. As $|N_G(a) \cap N_G(a') \cap D| \ge u^{1-\varepsilon/2}$ for all $\{a, a'\} \subset A$ or $\{a, a'\} \subset B$, we never run out of common neighbours in D. $\square$

We conclude this section by showing that any supposed counterexample to Theorem 1.1 can be partitioned almost entirely into quite large hubs.

Corollary 4.3. *Given $\varepsilon > 0$ there is $C \ge 1$ so that the following holds for all $\ell, n \in \mathbb{N}$ with $n \ge 3$ and $\ell \ge C\frac{\log n}{\log\log n}$. Suppose G is a C_ℓ-free graph on $N = (\ell-1)(n-1)+1$ vertices with $\alpha(G) \le n-1$. Then there is a partition $V(G) = W \cup \bigcup_{i\in[L]}(A_i \cup B_i \cup D_i)$ so that $|W| \le \varepsilon N$ and each (A_i, B_i, D_i) is a (u, ε)-hub with $u := \ell^{1-\varepsilon}$.*

Proof. Let $\delta > 0$ be such that Lemma 4.2 applies with $\varepsilon/2$ in place of ε. Let $\beta = \delta/7$ and $C \ge 4/\beta$ be sufficiently large. It suffices to show that any $W \subset V(G)$ with $|W| > \varepsilon N$ contains a (u, ε)-hub, as then iteratively removing such hubs proves the lemma.

To see this, we claim that we can apply Lemma 3.2 to $G[W]$ with $\gamma = \ell^\beta$, $D = \ell$ and $d = \ell^{1-\beta}$. Indeed, for C large we have $d \ge 8\gamma^2$ and $\alpha(G) \le n-1 \le |W|/d$, and also $D = \ell \ge \frac{3\log(N)}{\log(\ell^\beta)} \ge 3\log_\gamma(|W|)$, as $\ell \ge (4/\beta)\frac{\log n}{\log\log n}$. Thus Lemma 3.2 gives a subgraph H of $G[W]$ with $v(H) \le \ell$ and $e(G[U]) \ge d^2/2^9\gamma^4 = 2^{-9}\ell^{2-6\beta} \ge \ell^{2-\delta}$. Now Lemma 4.2 gives a (u, ε)-hub in $G[W]$. $\square$

5 Stability

In this section we upgrade the decomposition provided by Corollary 4.3 to obtain our main stability result, namely that any supposed counterexample to Theorem 1.1 can be partitioned almost entirely into quite large approximate cliques, and furthermore there are no edges between parts. The precise statement is as follows.

Lemma 5.1. *Given $\eta > 0$ there is $C \geq 1$ so that the following holds for all $\ell, n \in \mathbb{N}$ with $n \geq 3$ and $\ell \geq C \frac{\log n}{\log \log n}$. Suppose G is a C_ℓ-free graph on $N = (\ell - 1)(n - 1) + 1$ vertices with $\alpha(G) \leq n - 1$. Then there are disjoint sets $V_1, \dots, V_s \subset V(G)$ such that:*

(i) $|V_i| \in [(1 - \eta)\ell, \ell]$ *for all* $i \in [s]$*;*

(ii) $|\bigcup_{i \in [s]} V_i| \geq (1 - \eta)N$*;*

(iii) $G[V_i]$ *has minimum degree at least* $(1 - \eta)\ell$ *for all* $i \in [s]$*;*

(iv) There are no edges of G *between* V_i *and* V_j *for all distinct* $i, j \in [s]$*.*

Throughout the section we will fix G as in Lemma 5.1, with $\varepsilon < \varepsilon_0(\eta)$ sufficiently small and C sufficiently large so that Corollary 4.3 gives a partition $V(G) = W \cup \bigcup_{i \in [L]}(A_i \cup B_i \cup D_i)$ with $|W| \leq \varepsilon N$, where each (A_i, B_i, D_i) is a (u, ε)-hub with $u = \ell^{1-\varepsilon}$.

The proof proceeds in several stages, gradually refining the structure provided from the hubs to that in Lemma 5.1. In the next subsection we show how to find cycles of specified lengths in a system of hubs and 'handles' (suitable paths connecting the hubs). There is a potential parity obstacle due to the bipartite structure of hubs, but we can eliminate this obstacle using the bound on $\alpha(G)$; this is achieved in the second subsection. In the third subsection we study the interaction between hubs: roughly speaking, we consider an auxiliary graph H_3, where $V(H_3)$ consists of most of the hubs and we join two hubs if they are connected by a large matching. We show that H_3 cannot have large components, and then in the final subsection we show that these components identify the approximate cliques needed to prove Lemma 5.1.

5.1 Cycles from hubs and handles

In this subsection we show how to find cycles from a suitable system of hubs and connecting paths. Our first lemma concerns the following condition under which we can drop the parity restriction on lengths of paths within a hub. We say that a (u, ε)-hub (A, B, D) is *parity broken* if $G[A]$ contains a matching of size $2u^{1-\varepsilon}$.

Lemma 5.2. *Suppose (A, B, D) is a parity broken (u, ε)-hub in G. Let $s_1, \dots, s_m, t_1, \dots, t_m \in A \cup B$ be distinct and $\ell_1, \dots, \ell_m \geq 2$ with $\sum_{i \in [m]}(\ell_i + 1) \leq 2(1 - \varepsilon)u$. Suppose also that, for each $i \in [m]$, if ℓ_i is not a bipartite length for $\{s_i, t_i\}$ in $G[A, B]$ then $\ell_i \geq 7$. Then there are vertex-disjoint paths $P_1, \dots, P_m$ in $G[A \cup B \cup D]$, where each P_i is an $s_i t_i$-path of length ℓ_i.*

Proof. As (A, B, D) is parity broken and $2u^{1-\varepsilon} - 2m \geq m$, there is a matching $\mathcal{M} = \{x_i y_i : i \in [m]\}$ in $G[A]$ which is vertex-disjoint from $\{s_1, \dots, s_m, t_1, \dots, t_m\}$. We will apply the connection property of (A, B, D) to a collection of pairs $(s_{i,k}, t_{i,k})$ where there are one or two pairs for each original pair (s_i, t_i). If ℓ_i is a bipartite length for $\{s_i, t_i\}$ then we take one pair $(s_{i,1}, t_{i,1}) = (s_i, t_i)$ with the same length $\ell_{i,1} = \ell_i$. Otherwise, we take two pairs $(s_{i,1}, t_{i,1}) = (s_i, x_i)$ and $(s_{i,2}, t_{i,2}) = (y_i, t_i)$ with lengths $\ell_{i,1}, \ell_{i,2} \geq 2$ chosen such that both $\ell_{i,k}$ are bipartite lengths for $\{s_{i,k}, t_{i,k}\}$ in $G[A, B]$ with $\ell_{i,1} + \ell_{i,2} + 1 = \ell_i$. By the connection property of (A, B, D) we find vertex-disjoint $s_{i,k} t_{i,k}$-paths of lengths $\ell_{i,k}$, which combine with edges from $\mathcal{M}$ to produce the required paths $P_1, \dots, P_m$. □

Let $\mathcal{H}$ be a set of vertex-disjoint (u, ε)-hubs and $\mathcal{P} = \{P_1, \dots, P_k\}$ be a set of vertex-disjoint paths in a graph G. Suppose P_i is an $b_i a_{i+1}$-path for $i \in [k]$, writing $a_{k+1} := a_1$. We call $\mathcal{P}$ a *handle system* for $\mathcal{H}$ if

(i) each P_i is internally disjoint from $\bigcup\{V(H) : H \in \mathcal{H}\}$,

(ii) for each $i \in [k]$ there is $H_i \in \mathcal{H}$ with $\{a_i, b_i\} \subset V(H_i)$,

(iii) each $H \in \mathcal{H}$ contains at most $u^{1-\varepsilon}/2$ of $\{a_1, \dots, a_k\}$.

Note that we often apply the above definition with some paths P_i consisting only of the edge $b_i a_{i+1}$ (in which case condition (i) is vacuous). The next lemma shows how handle systems provide cycles of specified lengths.

Lemma 5.3. *Let $\mathcal{H}$ be a set of vertex-disjoint (u,ε)-hubs in G and $\mathcal{P} = \{P_1, \dots, P_k\}$ be a handle system for $\mathcal{H}$, where each P_i is an $b_i a_{i+1}$-path of length ℓ_i. Let $\ell_{\text{sum}} = \sum_{i\in[k]} \ell_i$. Then:*

(i) If $\{a_i, b_i\} \subset A_i$ for all $i \in [k]$ then G contains an ℓ-cycle for any $\ell \in [2k + \ell_{\text{sum}}, 2(1-\varepsilon)u|\mathcal{H}| + \ell_{\text{sum}} - 2k]$ of the same parity as ℓ_{sum}.

(ii) If some $\{a_j, b_j\}$ with $j \in [k]$ is contained in a parity broken hub of $\mathcal{H}$ then G contains an ℓ-cycle for any $\ell \in [7k + \ell_{\text{sum}}, 2(1-\varepsilon)u|\mathcal{H}| + \ell_{\text{sum}} - 2k]$.

Proof. We write $\ell - \ell_{\text{sum}} = \sum_{i\in[k]} \ell'_i$, where $\ell_j \geq 7$ (for (ii)), each $\ell'_i \geq 2$ with $i \neq j$ is a bipartite length for its hub, and for each $H \in \mathcal{H}$ we have $\sum\{\ell'_i + 1 : \{a_i, b_i\} \subset V(H)\} \leq 2(1-\varepsilon)u$. By the connection property of hubs, and Lemma 5.2 for the parity broken hub, we can find vertex-disjoint $a_i b_i$-paths of length ℓ'_i for each $i \in [k]$, which combine with $\mathcal{P}$ to produce an ℓ-cycle. □

5.2 Breaking parity

In this subsection we will prove that almost all hubs of G are parity broken. This will use the bound on the independence number of G, via the following proposition.

Proposition 5.4. *Let $m, d, s \in \mathbb{N}$ with $m \geq 3d$. Suppose G is a graph with $V(G) = \bigcup_{i\in[s]} I_i$, where $I_1, \dots, I_s$ are disjoint independent sets of order m. Suppose also that $\alpha(G) < v(G)/12d$. Then there is $\{i_0, \dots, i_d\} \subset [s]$ and a matching of size d with one edge in each $G[I_{i_{j-1}}, I_{i_j}]$ for $j \in [d]$.*

Proof. Consider a maximal matching $\mathcal{M}'$ in G with the property that $\mathcal{M}'$ contains at most one edge of $G[I_i, I_j]$ for all distinct $i, j \in [s]$. We use $\mathcal{M}'$ to define a graph H with $V(H) = [s]$, where $ij \in E(H)$ if and only if $\mathcal{M}'$ contains an edge from $G[I_i, I_j]$. To prove the proposition, it suffices to show that H contains a path of length d. By Theorem 2.5, it suffices to prove $d(H) > d - 1$.

For contradiction, suppose $d(H) \leq d - 1$. Let $S \subset V(H)$ with $|S| = s/2$ be such that $d_H(i) \leq d_H(j)$ for all $i \in S$, $j \notin S$. Then $d_H(i) \leq 2(d-1)$ for all $i \in S$. By Theorem 2.2 (Turán's Theorem), there is an independent set $S' \subset S$ in H with $|S'| \geq |S|/(2d-1) \geq s/4d$. For each $i \in S'$, let J_i be obtained from I_i by deleting all vertices contained in an edge of $\mathcal{M}'$. By the definition of $\mathcal{M}'$ and S, we have $|J_i| \geq |I_i| - 2d \geq m/3$. Since $\mathcal{M}'$ is maximal, there are no edges between J_i and J_j for any distinct i, j, so $\bigcup_{i\in S'} J_i$ is independent. We deduce $\alpha(G) \geq |S'|(m/3) \geq ms/12d = v(G)/12d$. This contradiction completes the proof. □

We can now show that almost all (u,ε)-hubs of G are parity broken.

Lemma 5.5. *At least $(1-\varepsilon)L$ hubs are parity broken.*

Proof. First we note that if $u \geq 4n$ then every hub (A, B, D) must be parity broken. Indeed, as $\alpha(G) < n$, any maximal matching in A has size at least $u/3 > 2u^{1-\varepsilon}$. Thus we may assume $n \geq u/4 = \ell^{1-\varepsilon}/4$.

For contradiction, suppose the hubs $\{(A_i, B_i, D_i)\}_{i\in[s]}$ are not parity broken, where $s = \varepsilon L \geq \varepsilon N/4u$. We will obtain a contradiction by using Lemma 5.3 to find an ℓ-cycle. Specifically, it suffices to show that there is a set of hubs $\mathcal{H} = \{H_1, \dots, H_k\}$ for some $k \geq \ell/u$, and a handle system $\mathcal{P} = \{P_1, \dots, P_k\}$ for $\mathcal{H}$, where each P_i has length ℓ_i, starts in H_i and ends in H_{i+1}, and $\ell_{\text{sum}} = \sum_{i\in[k]} \ell_i \leq \ell/4$ has the same parity as ℓ.

To achieve this, we look for a cycle of suitable length in the auxiliary graph H with $V(H) = [s]$, where $ij \in E(H)$ if and only if there is an edge between A_i and A_j. We apply Lemma 2.9 to H with $\gamma = \ell^{1-2\varepsilon}$ to obtain triples $\{(x_i, X_i, T_i)\}_{i\in[t]}$ so that for each $i \in [t]$ there is $d_i \in [0, \log_\gamma(N)]$ such that each vertex of X_i is at distance d_i from x_i in T_i. We let $X = \bigcup_{i\in[t]} X_i$ and note that $|X| \geq s/2\gamma \geq \varepsilon N/8u\gamma \geq \varepsilon n\ell^{-1+3\varepsilon}/16$.

We will construct a cycle by applying Proposition 5.4 to find a long path in $H[X]$. Consider a maximal matching in each $G[A_i]$ and let $I_i \subset A_i$ denote the vertices not covered by the matching. Then each I_i is independent and $|I_i| \geq u/2$ as (A_i, B_i, D_i) is not parity broken. Deleting some vertices if necessary we can assume $|I_i| = u/2$ for all $i \in [s]$. Fix $d \in \mathbb{N}$ of the same parity as ℓ with $d = \ell^\varepsilon + 2 \pm 1$. Then $u/2 > 3d$, and for large ℓ we have

$$\Big|\bigcup_{i\in X} I_i\Big|/12d \geq (\varepsilon n\ell^{-1+3\varepsilon}/16) \cdot (\ell^{1-\varepsilon}/24d) > n > \alpha(G).$$

Thus Proposition 5.4 applies to $G\big[\bigcup_{i\in X} I_i\big]$, giving some $\{i_0, \dots, i_d\} \subset X$ and a matching M of size d with one edge in each $G[I_{i_{j-1}}, I_{i_j}]$ for $j \in [d]$.

Note that $P = i_0 \dots i_d$ is a path in H, so Lemma 2.9 (iii) implies that it is contained in some X_i. By the distance property of X_i, the unique $i_0 i_d$-path Q in T_i is internally disjoint from P and has length $\ell(Q)$ which is even with $\ell(Q) \leq 2d_i \leq 2\log_\gamma(N)$. Let S be a set of edges obtained by choosing one edge in $G[A_x, A_y]$ for each edge xy of Q (which exists by definition of H). Then $M \cup S$ consists of a set of vertex-disjoint paths, which we denote $P_1, \dots, P_k$, with lengths $\ell_1, \dots, \ell_k$, where $k \geq d - 1 > \ell/u$ (as M is a matching) and $\ell_{\mathrm{sum}} = \sum_{i\in[k]} \ell_i = \ell(Q) + d \leq \ell/4$ has the same parity as ℓ. Furthermore, $\{P_1, \dots, P_k\}$ is a handle system for a set of hubs $\{H_1, \dots, H_k\}$ such that each P_i starts in H_i and ends in H_{i+1}. Now Lemma 5.3 (i) gives an ℓ-cycle, which is the required contradiction. $\square$

Remark: Henceforth, we will assume all hubs of G are parity broken. This can be guaranteed by taking ε slightly smaller in Corollary 4.3 and moving into W any hubs that are not parity broken.

5.3 Interaction between hubs

We will now organise most of the hubs into 'components', so that there is no large matching between two hubs in different components. To do so, we write $U_i = A_i \cup B_i$ for each $i \in [L]$ and consider a maximum matching $\mathcal{M}$ in $G[\cup_i U_i]$ such that (a) every $uv \in \mathcal{M}$ goes between distinct hubs, and (b) between any two distinct hubs there is at most one edge of $\mathcal{M}$. We define an auxiliary graph H_1 on $[L]$ where $ij \in E(H_1)$ iff there is an edge of $\mathcal{M}$ between U_i and U_j. We start by bounding the average degree of H_1.

Lemma 5.6. *H_1 has average degree at most $\ell^{1-3\varepsilon}$.*

Proof. For contradiction, suppose $d(H_1) \geq \ell^{1-3\varepsilon}$. We apply Lemma 2.8 to H_1 with $\gamma = \ell^{1-5\varepsilon}$ and $d_1 = \ell^\varepsilon$, noting that $d(H_1) \geq 16\gamma d_1$, to find an ℓ_1-cycle for some $\ell_1 \in [d_1, d_1 + 2\log_\gamma(N)] \subset [d_1, d_1 + \ell/10]$, using $\ell \geq C \log n / \log\log n$. Its edges correspond to a submatching $\mathcal{M}'$ of $\mathcal{M}$ of size ℓ, which forms a handle system for a set of ℓ hubs. As $8\ell_1 \leq \ell \leq u\ell_1$ and each hub is parity broken, Lemma 5.3 (ii) gives an ℓ-cycle, which is a contradiction. $\square$

By Lemma 5.6, at most εL vertices of H_1 have degree greater than $\varepsilon^{-1}\ell^{1-3\varepsilon}$ in H_1. Let H_2 be obtained from H_1 by deleting these high degree vertices, so that $v(H_2) \geq (1-\varepsilon)L$. We will now restrict attention to the subgraph H_3 of H_2 where $ij \in E(H_3)$ iff $G[U_i, U_j]$ has a matching of size $2\ell^\varepsilon$. We show that H_3 does not have large components.

Lemma 5.7. *All connected components of H_3 have fewer than $(1+2\varepsilon)\ell^\varepsilon/2$ vertices.*

Proof. For contradiction, suppose H_3 contains a tree T with $(1+2\varepsilon)\ell^\varepsilon/2$ vertices. By definition of H_3, we can greedily choose a matching $\mathcal{P} = \{P_1, \dots, P_k\}$ that contains two edges of $G[U_i, U_j]$ for each $ij \in E(T)$. We can regard $\mathcal{P}$ as a handle system for the hubs $\mathcal{H} = \{(A_i, B_i, D_i) : i \in V(T)\}$. To see this, we note that condition (i) is vacuous, and (iii) holds as $|\mathcal{P}| = 2e(T) < \ell^\varepsilon < u^{1-\varepsilon}/2$. To achieve (ii), we order the edges of $\mathcal{P}$ cyclically according to a closed walk in T that uses every edge exactly twice (which is well-known to exist, e.g. by embedding T in the plane and walking around its outside). As $8|\mathcal{P}| \le \ell \le 2(1-\varepsilon)u|\mathcal{P}| - |\mathcal{P}|$ and all hubs are parity broken, Lemma 5.3 (ii) gives an ℓ-cycle, which is a contradiction. □

5.4 Proof of stability

We now combine the results of this section to prove our stability result.

Proof of Lemma 5.1. Let the graphs H_1, H_2 and H_3 be as in the previous subsection. Fix a maximal matching M_{ij} in $G[U_i, U_j]$ for each $ij \in E(H_2) \setminus E(H_3)$; by definition of H_3 each $|M_{ij}| \le 2\ell^\varepsilon$. For each $i \in V(H_2) = V(H_3)$, let $U'_i = U_i \setminus \bigcup_{ij} V(M_{ij})$; by definition of H_2 each $|U'_i| \ge |U_i| - \varepsilon^{-1}\ell^{1-3\varepsilon} \cdot 2\ell^\varepsilon \ge (1-\varepsilon)2u$ for large ℓ. Let $U' = \bigcup\{U'_i : i \in V(H_3)\}$ and $G' = G[U']$. We have $|U'| \ge |V(H_3)| \cdot (1-\varepsilon)2u \ge (1-\varepsilon)^2 2uL \ge (1-3\varepsilon)N$, so by Theorem 2.2 (Turán's Theorem) $d(G') \ge (1-4\varepsilon)\ell$.

Note that all edges of G' lie within some hub or join two hubs in the same connected component of H_3. By Lemma 5.7 the number of vertices in any component of G' is at most $(1+2\varepsilon)(\ell^\varepsilon/2) \cdot (2\ell^{1-\varepsilon}) = (1+2\varepsilon)\ell$. Let B be obtained from U' by deleting $V(C)$ for any component C of G' with $d(C) \le (1-\varepsilon^{1/2})\ell$. Then $|U'|(1-4\varepsilon)\ell \le 2e(G') \le |B|(1+\varepsilon)\ell + (|U'|-|B|)(1-\varepsilon^{1/2})\ell$, which gives $|B|(\varepsilon^{1/2}+\varepsilon) \ge |U'|(\varepsilon^{1/2} - 4\varepsilon)\ell$, and so $|B| \ge (1-6\varepsilon^{1/2})|U'| \ge (1-7\varepsilon^{1/2})N$.

We conclude by taking subgraphs of high minimum degree in each component of $G'[B]$. Letting $k = (1-\eta/2)\ell$, each such component C has $e(C) = d(C)v(C)/2 \ge (1-\varepsilon^{1/2})\ell v(C)/2 \ge \binom{k}{2} + (v(C)-k)(1-\eta/2)\ell$, as $(1+\varepsilon)\ell \ge v(C) \ge d(C) \ge (1-\varepsilon^{1/2})\ell$ and $\varepsilon \ll \eta$. Proposition 2.1 (i) gives a subgraph C' of C with $\delta(C') \ge k \ge (1-\eta/2)\ell \ge (1-3\eta/4)v(C)$. We let $V_1, \dots, V_s$ be the vertex-sets of these subgraphs C' for all components C of $G'[B]$. Then each $|V_i| \ge \delta(G[V_i]) \ge (1-\eta)\ell$ and $\sum_{i=1}^s |V_i| \ge (1-3\eta/4)|B| \ge (1-\eta)N$. Lastly, suppose for contradiction that some $|V_i| \ge \ell$. We may delete $|V_i| - \ell \le 3\varepsilon\ell$ vertices from V_i and apply Theorem 2.3 (Dirac's Theorem) to find an ℓ-cycle in G. This contradiction shows that all $|V_i| \le \ell - 1$. □

6 The upper bound

In this section we will prove our main result, Theorem 1.1, which establishes the upper bound on cycle-complete Ramsey numbers; the proof will be given in the last subsection. Most of this section will be occupied with cleaning up the approximate structure of a supposed counterexample, as provided by the stability result in the last section, until it becomes clear that its properties are contradictory, so it cannot exist.

Throughout the section we fix a graph G and 'approximate cliques' $V_1, \dots, V_s$ satisfying the hypotheses and conclusions of Lemma 5.1. In the first subsection we give conditions under which the approximate cliques can absorb additional vertices from the remainder $R := V(G) \setminus \bigcup_{i=1}^s V_i$, while maintaining pancyclicity and also the property that any pair of vertices can be connected by paths with a large range of possible lengths. In the second subsection we clean up R by absorbing some of its vertices into the approximate cliques. In the third subsection we show that the remaining part of R can be separated from most of the approximate cliques, in the sense they have each have a large subset with no neighbours in R. In the fourth subsection we show that one of the approximate cliques has a vertex that can absorb its neighbours. This final property quickly leads to a contradiction, which will complete the proof.

6.1 Absorbable paths

In this subsection we consider the following set-up which is very similar to the handle systems used for hubs. Given a set of paths $\mathcal{P} = \{P_1, \dots, P_m\}$ in a graph H and a set $V \subset V(H)$, we say $\mathcal{P}$ is absorbable into V if it consists of paths that are vertex-disjoint and disjoint from V, and there are distinct vertices $\{a_1, \dots, a_m, b_1, \dots b_m\} \subset V$ such that a_i is adjacent to one end of P_i and b_i is adjacent to the other end of P_i; we say that P_i attaches to a_i and b_i. The following lemma will be used to absorb paths into approximate cliques.

Lemma 6.1. *Let H be a graph with a partition $V(H) = U \cup V$, where $\delta(H[V]) \geq 0.9|V|$ and $|U| \leq 0.1|V|$. Suppose that $\mathcal{P}$ is a set of paths of length at most 2 which is absorbable into V and has $\cup_{P \in \mathcal{P}} V(P) = U$. Then*

(i) H contains an xy-path of length ℓ for any distinct x, y in $V(H)$ and $\ell \in [6, 2v(H)/3]$,

(ii) H is pancyclic.

Proof. For (i), we suppose first that both x and y are in V, and show that there is an xy-path of length ℓ for any $\ell \in [2, 2v(H)/3]$. To see this, we use $\delta(H[V]) \geq 0.9|V|$ to greedily choose an xy'-path P of length $\ell-2$ in $H[V]$ that avoids y. As $|N_{H[V]}(y) \cap N_{H[V]}(y')| - |V(P)| \geq 0.8|V| - (2/3)1.1|V| > 0$ we can choose a common neighbour of y and y' in $V \setminus V(P)$, and so obtain the required xy-path of length ℓ. Next we suppose that x is in V and y is in U. Then y lies on a path $P \in \mathcal{P}$. Let a and b be the attachments of P, where without loss of generality $a \neq x$. The subpath of P from y to a has length $\ell' \leq 3$. Adding a path of length $\ell - \ell'$ from a to x gives the required xy-path of length ℓ. Finally, suppose x and y are both in U. Then we can find a and b in V so that there is an xa-path and yb-path that are vertex-disjoint and both of length at most 2. Adding an ab-path of the appropriate length completes the proof of (i).

For (ii), we first note that by Theorem 2.4 (Bondy's Theorem) $H[V]$ is pancyclic. It remains to show there is an ℓ-cycle whenever $|V| < \ell \leq |V(H)|$. Let S be the set of attachments of $\mathcal{P}$, and fix any $V' \subset V \setminus S$ with $|V'| = \ell - |S| - |U|$. As $|U| + |S| \leq 3|U| \leq 0.3|V|$, we have $|V'| \geq 0.7|V|$. Let H' be the graph obtained from $H[V']$ by adding a new vertex v_P for each $P \in \mathcal{P}$, which is joined to all common neighbours in V' of the attachments of P. Note that $v(H') = |V'| + |\mathcal{P}|$ and $\delta(H') \geq |V'| - 0.2|V| \geq |V|/2 > v(H')/2$, and so by Theorem 2.3 (Dirac's Theorem) H' has a Hamilton cycle. Replacing each v_P by P and the edges to its attachments produces a cycle of length ℓ in H, as required. □

6.2 Cleaning up the remainder

Here we clean up the remainder $R = V(G) \setminus \bigcup_{i=1}^{s} V_i$ by absorbing some of its vertices into the approximate cliques, according to the following algorithm. For each $i \in [s]$ we keep track of two sets during the algorithm: (a) a set $W_i = V_i \cup R_i$, where $R_i \subset R$ has been absorbed by V_i, and (b) a subset A_i of V_i, which is available for further attachments in the sense of the previous subsection. We start with $W_i = A_i = V_i$ for each $i \in [s]$. In a given round:

- Consider any path P of length at most 2 in $G[R]$ that attaches to some distinct vertices a, b in A_i for some $i \in [s]$. If there is no such P then stop. Otherwise, move $V(P)$ from R to R_i, delete a and b from A_i, and proceed to the next round.

We claim that the algorithm terminates with $|W_i| \le \ell - 1$ for all $i \in [s]$. Indeed, otherwise in some round some $|W_i| \in [\ell, \ell+2]$, as W_i increments by at most 3 vertices in each round. Then W_i has a partition $W_i = V_i \cup R_i$, where $|V_i| \ge \delta(G[V_i]) \ge (1-\eta)|V_i| \ge 0.9|V_i|$ and $|R_i| \le \eta\ell + 2 \le 0.1|V_i|$. By construction, R_i is the union of paths $\mathcal{P}_i$ absorbable into V_i, so Lemma 6.1 (ii) gives an ℓ-cycle in $G[W_i]$. This contradiction proves the claim. We deduce $|R_i| = |W_i| - |V_i| < \eta\ell$. Furthermore, each A_i decreased by two vertices for each path added to R_i, so $|A_i| \ge |V_i| - 2|R_i| \ge (1-3\eta)\ell$.

6.3 Separating the remainder

Now we show that the cleaned up remainder $R := V(G) \setminus \bigcup_{i=1}^s W_i$ can be separated from most of the approximate cliques, in the following sense. For $i \in [s]$ let A'_i be the set of $v \in A_i$ such that v has a neighbour in R. We partition $[s]$ as $S \cup T$, where $T = \{i \in [s] : |A'_i| < \ell^{2/3}\}$.

Lemma 6.2. $|T| \ge s/2$.

Proof. We start by constructing a partition $R = U_1 \cup \cdots \cup U_r$, where each $G[U_j]$ has diameter at most 2 and $r \le 2N\ell^{-1/2}$. To see that this is possible, we repeatedly remove stars from R of order $\ell^{1/2}$ until none remain. We can remove at most $\eta N/\ell^{1/2}$ such stars. The remaining set R' must have $d(G[R']) < \ell^{1/2} - 1$. By Theorem 2.2 (Turán's Theorem) $|R'|/\ell^{1/2} \le \alpha(G[R']) \le \alpha(G) < n$, so $|R'| < n\ell^{1/2} < \frac{3}{2}N\ell^{-1/2}$. We let the parts $U_1, \dots, U_r$ consist of all removed stars and singleton parts for each vertex of R'. Then $r \le 2N\ell^{-1/2}$, as required.

Now suppose for contradiction that $|T| < s/2$, so $|S| > s/2$. For each $v \in \bigcup_{i=1}^s A'_i$ we fix any $u_v \in N(v) \cap R$. We consider an auxiliary bipartite graph H with parts $A = \{W_i\}_{i \in [s]}$ and $B = \{U_j\}_{j \in [r]}$, where we add an edge from W_i to U_j for each $v \in A'_i$ with $u_v \in U_j$. To see that this gives a (simple) graph we use the termination condition of the algorithm in the previous subsection: there cannot be distinct $v_1, v_2 \in A'_i$ with neighbours $u_1, u_2 \in U_j$, as U_j has diameter at most 2, so we would have a u_1u_2-path of length at most 2 attaching to A_i.

We will obtain a contradiction by finding a short cycle in H and using it to construct an ℓ-cycle in G. We have $v(H) = s + r \le 2N\ell^{-1} + 2N\ell^{-1/2} \le 4N\ell^{-1/2}$ and $e(H) = \sum_{i \in [s]} |A'_i| \ge |S|\ell^{2/3} > \frac{1}{2}\frac{N}{2\ell}\ell^{2/3} \ge \ell^{1/6}v(H)/16$, so $d(H) > \ell^{1/6}/8$. As $\ell \ge C\frac{\log n}{\log\log n}$, we can apply Lemma 2.8 with $\gamma = d_1 = \ell^{1/14}$ to find a cycle in H with length in $[\ell^{1/14}, \ell^{1/14} + 2\log_{\ell^{1/14}}(v(H))] \subset [4, \ell/8]$.

As H is bipartite, we can write this cycle as $W_{i_1}U_{i_1}W_{i_2}\cdots W_{i_L}U_{i_L}W_{i_1}$, for some $2 \le L \le \ell/16$. Each U_{i_j} has diameter at most 2, so by construction of H there is a path Q_j of length at most 4, starting with the edge $b_j u_{b_j}$ for some $b_j \in W_{i_j}$ and ending with the edge $u_{a_{j+1}}a_{j+1}$ for some $a_{j+1} \in W_{i_{j+1}}$. Furthermore, $a_j, b_j \in W_{i_j}$ are distinct, as $u_{a_j} \ne u_{b_j}$. We fix $\ell_j \in [2, \ell/2]$ for each $j \in [L]$ with $\sum_{j\in[L]} \ell_j = \ell - \sum_{j \in [L]} e(Q_j)$. and apply Lemma 6.1 (i) to choose a_jb_j-paths P_j in W_j of length ℓ_j. Combining these with the paths Q_j produces an ℓ-cycle, which is a contradiction. □

6.4 Absorbing neighbours

Now we will show that one of the approximate cliques has a vertex that can absorb its neighbours. To do so, we now analyse the edges crossing between the approximate cliques. For each $i \in T$ let $B_i = A_i \setminus A'_i$ denote the set of $v \in A_i$ with no neighbour in R. By definition of T each $|B_i| \ge |A_i| - \ell^{2/3} \ge 2|V_i|/3$. For each $i \in T$ we consider a matching $\mathcal{M}_i$ in G of maximum size subject to the condition that each edge of $\mathcal{M}_i$ intersects W_i in a single vertex from B_i. We will show that these matchings cannot all be large.

Lemma 6.3. *There is $i^* \in T$ with $|\mathcal{M}_{i^*}| \le \ell^{1/3}$.*

Before giving the proof, we show how this lemma allows us to find a vertex that can absorb its neighbours. Recall that $W_{i^*} = V_{i^*} \cup R_{i^*}$ and R_{i^*} is a union of vertex-disjoint paths $\mathcal{P}_{i^*}$ that is absorbable into V_{i^*}.

Lemma 6.4. *There is $v \in B_{i^*}$ such that for any neighbours $y_1, \dots, y_k$ of v in $\overline{W_{i^*}}$ with $k \le \ell - |W_{i^*}|$, letting $\mathcal{P}'_{i^*}$ be obtained from $\mathcal{P}_{i^*}$ by adding each y_i as a path of length 0, we have $\mathcal{P}'_{i^*}$ absorbable into V_{i^*}.*

Proof. We apply the following algorithm to construct a set $D_{i^*} \subset B_{i^*}$ such that every vertex in B_{i^*} has the stated property. We start with $D_{i^*} = B_{i^*}$ and $X = \emptyset$. While there is $x \in \bigcup_{j \ne i^*} W_j$ with $1 \le d_G(x, D_{i^*}) \le 2\ell^{1/3}$ we add x to X and delete $N_G(x) \cap D_{i^*}$ from D_{i^*}. This process terminates with a set D_{i^*} such that $d_G(x, D_{i^*}) = 0$ or $d_G(x, D_{i^*}) > 2\ell^{1/3}$ for all $x \in (\bigcup_{j \ne i} W_j) \setminus X$. Each $x \in X$ has a private neighbour in B_{i^*}, so by choice of $\mathcal{M}_{i^*}$ we have $|X| \le |\mathcal{M}_{i^*}| \le \ell^{1/3}$, and so $|D_{i^*}| \ge |B_{i^*}| - (2\ell^{1/3})|X| \ge \ell/2 > 0$.

Consider any $v \in D_{i^*}$ and neighbours $y_1, \dots, y_k$ of v in $\overline{W_{i^*}}$ with $k \le \ell - |W_{i^*}|$. Each y_i is not in X (otherwise we would have deleted v from D_{i^*}) so has at least $2\ell^{1/3}$ neighbours in D_{i^*}. This implies $k \le |\mathcal{M}_{i^*}| \le \ell^{1/3}$, or otherwise we could greedily construct a matching of size $|\mathcal{M}_{i^*}| + 1$ between $\{y_1, \dots, y_k\}$ and B_{i^*}, which is contrary to the choice of $\mathcal{M}_{i^*}$. We can therefore greedily choose two attachments for each y_i in D_{i^*}, which are distinct from each other, and distinct from the attachments of $\mathcal{P}_{i^*}$ as $D_{i^*} \subset B_{i^*} \subset A_{i^*}$. Thus $\mathcal{P}'_{i^*}$ is absorbable into V_{i^*}. $\square$

We conclude this subsection by returning to the proof of Lemma 6.3.

Proof of Lemma 6.3. For contradiction, suppose $|\mathcal{M}_i| > \ell^{1/3}$ for all $i \in T$. Note that every edge in $\mathcal{M}_i$ has one end in B_i and the other end in $\cup_{j \ne i} W_j$ (it is not in R by definition of B_i). Consider a uniformly random partition $[s] = S_1 \cup S_2$. Say that $bc \in \mathcal{M}_i$ with $b \in B_i$ and $c \in W_j$ is *good* if $i \in S_1$ and $j \in S_2$. Each edge is good with probability $1/4$, so we can fix a partition so that the number of good edges is at least $\frac{1}{4} \sum_{i \in T} |\mathcal{M}_i| > |T|\ell^{1/3}/4 \ge s\ell^{1/3}/8$.

Consider the auxiliary bipartite graph H with parts $A = \{W_i\}_{i \in S_1}$ and $B = \{W_j\}_{j \in S_2}$, where we add an edge from $W_i \in A$ to $W_j \in B$ for each good edge $bc \in \mathcal{M}_i$ with $b \in B_i$ and $c \in W_j$. We claim that H is a (simple) graph. To see this, suppose on the contrary we have b_1c_1 and b_2c_2 in $\mathcal{M}_i$ with $\{c_1, c_2\} \subset W_j$. By Lemma 6.1 (i) there is a b_1b_2-path P_1 in $G[W_i]$ of length $\lfloor \ell/2 \rfloor - 1$ and a c_1c_2-path P_2 in $G[W_j]$ of length $\lceil \ell/2 \rceil - 1$. Combining the paths P_1 and P_2 with the edges b_1c_1 and b_2c_2 gives a ℓ-cycle. This contradiction proves the claim.

We deduce $e(H) \ge s\ell^{1/3}/8 = v(H)\ell^{1/3}/8$, so $d(H) \ge \ell^{1/3}/4$. We use this to obtain the required contradiction by finding a short cycle in H, and so an ℓ-cycle in G. This part of the proof is very similar to that of Lemma 6.2. Lemma 2.8 provides an even cycle $W_{i_1}W_{j_1} \cdots W_{i_L}W_{j_L}W_{i_1}$, for some $2 \le L \le \ell/16$, where each $i_\alpha \in S_1$ and $j_\alpha \in S_2$. By definition of H, for each $\alpha \in [L]$ there are edges $a_\alpha x_\alpha$ and $b_\alpha y_\alpha$ in $\mathcal{M}_{i_\alpha}$ with $\{a_\alpha, b_\alpha\} \subset W_{i_\alpha}$, $x_\alpha \in W_{j_\alpha - 1}$ and $y_\alpha \in W_{j_\alpha}$. By Lemma 6.1 (i) there is a path Q_α of length at most 4 from $b_{\alpha-1}$ to a_α through $W_{j_\alpha - 1}$ via $y_{\alpha-1}$ and x_α (whether or not these coincide). We fix $\ell_\alpha \in [2, \ell/2]$ for each $\alpha \in [L]$ with $\sum_{\alpha \in [L]} \ell_\alpha = \ell - \sum_{\alpha \in [L]} e(Q_\alpha)$. and apply Lemma 6.1 (i) to choose $a_\alpha b_\alpha$-paths P_α in W_{i_α} of length ℓ_α. Combining these with the paths Q_α produces an ℓ-cycle, and so the required contradiction. $\square$

6.5 Proof of Theorem 1.1

We now complete the proof of our main theorem.

Proof of Theorem 1.1. We fix $\ell \in \mathbb{N}$ and prove the following statement (*) by induction on $n \geq 1$ such that if $n \geq 3$ we have $\ell \geq C\frac{\log n}{\log\log n}$ (for some large absolute constant C):

(*) there is no C_ℓ-free graph G with $v(G) = N = (\ell-1)(n-1)+1$ and $\alpha(G) \leq n-1$.

The case $n = 1$ holds as every graph G with $v(G) \geq 1$ has an independent set of order 1. The case $n = 2$ holds as every graph G with $v(G) \geq \ell$ contains an independent set of order 2 or a clique of order ℓ.

Now we give the induction step for $n \geq 3$. For contradiction, suppose we have a C_ℓ-free graph G with $v(G) = N = (\ell-1)(n-1)+1$ and $\alpha(G) \leq n-1$.

If there is any vertex v of degree less than $\ell-1$ we delete $N(v)\cup\{v\}$ from G and apply induction. The remaining subgraph G_1 satisfies $v(G_1) \geq (\ell-1)(n-2)+1 = r(C_\ell, K_{n-1})$ by induction, so it contains a cycle C of length ℓ or an independent set I of order $n-1$. Then G contains an ℓ-cycle C or $\{v\}\cup I$ forms an independent set of order n. Thus we may assume $\delta(G) \geq \ell-1$.

We let $V_1, \dots, V_s$ be the approximate cliques provided by the stability result (Lemma 5.1), let $W_i = V_i \cup R_i$ for $i \in [s]$ be the enlarged approximate cliques obtained in the previous section by absorbing part of the remainder, and let $v \in B_{i^*}$ be given by Lemma 6.4. As v has at least $\ell-1$ neighbours, we can choose neighbours $y_1, \dots, y_k$ of v in $\overline{W_{i^*}}$ with $k = \ell - |W_{i^*}|$. As the path system $\mathcal{P}'_{i^*}$ in Lemma 6.4 is absorbable, Lemma 6.1 gives a cycle of length $|V_i| + |R_i| + k = \ell$ in G. This gives a contradiction and completes the proof of the theorem. $\square$

7 Concluding remarks

Our results answer the questions of Erdős et al. [21] up to a constant factor, which we did not compute explicitly, although with more work it seems that a reasonable value (less than 20, say) can be obtained. It would be interesting to obtain an asymptotic formula for the ℓ minimising $r(C_\ell, K_n)$. The constructions for the lower bound on $r(C_\ell, K_n)$ avoid a range of cycles. For large ℓ, this range consists of all cycles of length at least ℓ, and for small ℓ, it consists of all cycles of length at most ℓ. This suggests that the finer nature of the threshold may be connected to the problem of improving the Moore bound (see [39]) on the number of edges in a graph of given order and diameter.

The problem of obtaining good estimates on $r(C_\ell, K_n)$ for small $\ell > 3$ remains widely open. The most significant gap in the current state of knowledge is the case $\ell = 4$, for which the known bounds (see [12, 47]) are $c(n/\log n)^{3/2} \leq r(C_4, K_n) \leq C(n/\log n)^2$ for some constants c and C.

Acknowledgement. The authors would like to thank the organisers of the Workshop on Extremal and Structural Combinatorics held at IMPA in Rio de Janeiro, where this work began.

References

[1] P. Allen, G. Brightwell, J. Skokan, Ramsey-goodness – and otherwise, *Combinatorica* 33:125–160, 2013.

[2] N. Alon, Subdivided graphs have linear Ramsey numbers, *J. Graph Theory* 18:343–347, 1994.

[3] N. Alon and J. Spencer, *The probabilistic method* 4th ed., Wiley, 2016.

[4] M. Ajtai, J. Komlós and E. Szemerédi, A note on Ramsey numbers, *J. Combin. Theory Ser. A* 29:354–360, 1980.

[5] T. Bohman and P. Keevash, Dynamic concentration of the triangle-free process, *arXiv:1302.5963*.

[6] J.A. Bondy and P. Erdős, Ramsey numbers for cycles in graphs, *J. Combin. Theory Ser. B* 14:46–54, 1973.

[7] B. Bollobás, *Modern graph theory*, Springer-Verlag, 1998.

[8] B. Bollobás, C.J. Jayawardene, Z.K. Min, C.C. Rousseau, H.Y. Ru, and J. Yang, On a conjecture involving cycle-complete graph Ramsey numbers, *Australas. J. Combin.* 22:63–72, 2000.

[9] J.A. Bondy, Pancyclic graphs I, *J. Combin. Theory Ser. B* 11:80–84, 1971.

[10] S.A. Burr and P. Erdős, On the magnitude of generalized Ramsey numbers for graphs, in: Infinite and Finite Sets I (Colloq., Keszthely, 1973), Colloq. Math. Soc. Janos Bolyai 10:214–240, 1975.

[11] S.A. Burr and P. Erdős, Generalizations of a Ramsey-theoretic result of Chvátal, *J. Graph Theory* 7:39–51, 1983.

[12] Y. Caro, Y. Li, C. Rousseau and Y. Zhang, Asymptotic bounds for some bipartite graph-complete graph Ramsey numbers, *Disc. Math.* 220:51–56, 2000.

[13] G. Chen and R.H. Schelp, Graphs with linearly bounded Ramsey numbers, *J. Combin. Theory Ser. B* 57:138–149, 1993.

[14] F. Chung and R. Graham, *Erdős on graphs: his legacy of unsolved problems*, A.K. Peters, 1998.

[15] V. Chvátal and F. Harary, Generalized Ramsey theory for graphs, III. Small off-diagonal numbers, *Pacific J. Math* 41:335–345, 1972.

[16] V. Chvátal, V. Rödl, E. Szemerédi and W.T. Trotter Jr., The Ramsey number of a graph with bounded maximum degree, *J. Combin. Theory Ser. B* 34:239–243, 1983.

[17] D. Conlon, A new upper bound for diagonal Ramsey numbers, *Ann. of Math.* 170:941–960, 2009.

[18] G. Dirac, Some theorems on abstract graphs, *Proc. London Math. Soc.* 2:69–81, 1952.

[19] P. Erdős, Remarks on a theorem of Ramsey, *Bull. Res. Council Israel Sect. F* 7F:21–24, 1957.

[20] P. Erdős, Some remarks on the theory of graphs, *Bull. Amer. Math. Soc.* 53:292–294, 1947.

[21] P. Erdős, R.J. Faudree, C.C. Rousseau and R.H. Schelp, On cycle-complete graph Ramsey numbers, *J. Graph Theory* 2:53–64, 1978.

[22] P. Erdős and T. Gallai, On maximal paths and circuits of graphs, *Acta Math. Acad. Sci. Hungar.* 10:337–356, 1959.

[23] P. Erdős and G. Szekeres, A combinatorial problem in geometry, *Compositio Math.* 2:463–470, 1935.

[24] R.J. Faudree and R.H. Schelp, All Ramsey numbers for cycles in graphs, *Disc. Math* 8:313–329, 1974.

[25] G. Fiz Pontiveros, S. Griffiths and R. Morris, The triangle-free process and the Ramsey number $R(3,k)$, *Mem. Amer. Math. Soc*, to appear.

[26] J. Fox and B. Sudakov, Density theorems for bipartite graphs and related Ramsey-type results, *Combinatorica* 29:153–196, 2009.

[27] J. Fox and B. Sudakov, Two remarks on the Burr-Erdős conjecture, *Europ. J. Combin.* 30:1630–1645, 2009.

[28] J. Fox and B. Sudakov, Dependent random choice, *Random Struct. Alg.* 38:1–32, 2011.

[29] R. Graham, B. Rothschild and J. Spencer, *Ramsey Theory*, John Wiley & Sons, 1990.

[30] P. Keevash and B. Sudakov, Pancyclicity of Hamiltonian and highly connected graphs, *J. Combin. Theory Ser. B* 100:456–467, 2010.

[31] J.H. Kim, The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$, *Random Struct. Alg.* 7:173–207, 1995.

[32] J. Komlós and M. Simonovits, Szemerédi's regularity lemma and its applications in graph theory, in: Combinatorics, Paul Erdős is eighty, Bolyai Soc. Math. Stud. 2:295–352, János Bolyai Math. Soc., Budapest, 1996.

[33] A. Kostochka and V. Rödl, On graphs with small Ramsey numbers, *J. Graph Theory* 37:198–204, 2001.

[34] A. Kostochka and V. Rödl, On graphs with small Ramsey numbers II, *Combinatorica* 24:389–401, 2004.

[35] A. Kostochka and B. Sudakov, On Ramsey numbers of sparse graphs, *Combin. Prob. Comput.* 12:627–641, 2003.

[36] C. Lee, Ramsey numbers of degenerate graphs, *Ann. of Math.* 185:791–829, 2017.

[37] E. Long, Long paths and cycles in subgraphs of the cube, *Combinatorica* 33:395–428, 2013.

[38] E. Long, Long paths in the cube and other combinatorial results, PhD thesis, University of Cambridge, 2013.

[39] M. Miller and J. Širáň, Moore graphs and beyond: A survey of the degree/diameter problem, *Electron. J. Combin.* Dynamic survey DS14.

[40] V. Nikiforov, The cycle-complete graph Ramsey numbers, *Combin. Probab. Comput.* 14:349–370, 2005.

[41] A. Pokrovskiy and B. Sudakov, Ramsey goodness of paths, *J. Combin. Theory Ser. B* 122:384–390, 2017.

[42] F.P. Ramsey, On a Problem of Formal Logic, *Proc. London Math. Soc.* 30:264–286, 1930.

[43] V. Rosta, On a Ramsey type problem of J.A. Bondy and P. Erdős, I & II, *J. Combin. Theory Ser. B* 15:94–120, 1973.

[44] I. Schiermeyer, The Cycle-Complete Graph Ramsey Number $r(C_5, K_7)$, *Disc. Math. Graph Theory* 25:129–139, 2005.

[45] I. Schiermeyer, All cycle-complete graph Ramsey numbers $r(C_m, K_6)$, *J. Graph Theory* 44:251–260, 2003.

[46] J. Shearer, A note on the independence number of triangle-free graphs, *Discrete Math.* 46:83–87, 1983.

[47] J. Spencer, Asymptotic lower bounds for Ramsey functions, *Disc. Math.* 20:69–76, 1977.

[48] J. Spencer, Ramsey's theorem – a new lower bound, *J. Combin. Theory, Ser. A* 18:108–115, 1975.

[49] B. Sudakov, A note on odd cycle-complete graph Ramsey numbers, *Electron. J. Combin.* 9:N1, 2002.

[50] P. Turán, Eine Extremalaufgabe aus der Graphentheorie, *Mat. Fiz. Lapok.* 48:436–452, 1941.

[51] A. Thomason, An upper bound for some Ramsey numbers, *J. Graph Theory* 12:509–517, 1988.

[52] J.S. Yang, Y.R. Huang and K.M. Zhang, The value of the Ramsey number $R(C_n, K_4)$ is $3(n-1)+1$ $(n \geq 4)$, *Australas. J. Combin.* 20:205–206, 1999.