Strengthening Rödl's theorem

Maria Chudnovsky[1]
Princeton University, Princeton, NJ 08544

Alex Scott[2]
Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

Paul Seymour[3]
Princeton University, Princeton, NJ 08544

Sophie Spirkl[4]
University of Waterloo, Waterloo, Ontario N2L3G1, Canada

Abstract

What can be said about the structure of graphs that do not contain an induced copy of some graph H? Rödl showed in the 1980s that every H-free graph has large parts that are very sparse or very dense. More precisely, let us say that a graph F on n vertices is *ε-restricted* if either F or its complement has maximum degree at most εn. Rödl proved that for every graph H, and every $\varepsilon > 0$, every H-free graph G has a linear-sized set of vertices inducing an ε-restricted graph. We strengthen Rödl's result as follows: for every graph H, and all $\varepsilon > 0$, every H-free graph can be partitioned into a bounded number of subsets inducing ε-restricted graphs.

1   Introduction

What can be said about the structure of graphs that do not contain an induced copy of some graph H? In the 1980s, Rödl [7] showed that every H-free graph has large parts that are very sparse or very dense.

To say that more precisely we need some definitions. Graphs in this paper are finite and without loops or parallel edges. If G is a graph and $X \subseteq V(G)$, we denote the subgraph of G induced on X by $G[X]$, and $\overline{G}$ denotes the complement graph of G. If G, H are graphs, we say that G is H-free if no induced subgraph of G is isomorphic to H. For a graph G, let us say $X \subseteq V(G)$ is *weakly ε-restricted* if one of $G[X], \overline{G}[X]$ has at most $\varepsilon|X|^2$ edges; and that X is *ε-restricted* if one of the graphs $G[X], \overline{G}[X]$ has maximum degree at most $\varepsilon|X|$. Rödl [7] proved the following:

1.1 Theorem. *For every graph H, and all $\varepsilon > 0$, there exists $\delta > 0$ such that for every H-free graph G, there is a weakly ε-restricted set $X \subseteq V(G)$ with $|X| \geq \delta|G|$.*

Every ε-restricted set is weakly $\varepsilon/2$-restricted, and every weakly $\varepsilon/2$-restricted set has a subset of at least half its size that is ε-restricted. Thus, an equivalent version of Rödl's theorem is the following:

1.2 Theorem. *For every graph H, and all $\varepsilon > 0$, there exists $\delta > 0$ such that for every H-free graph G, there is an ε-restricted set $X \subseteq V(G)$ with $|X| \geq \delta|G|$.*

Rödl's theorem is an easy consequence of Szemerédi's regularity lemma, and has proved extremely useful. For example, it is now a standard tool in approaching the Erdős-Hajnal conjecture (see for instance the breakthrough paper [2], where it was crucial, and much subsequent work). A proof of 1.2 not using the regularity lemma (and consequently with much better constants) was given by Fox and Sudakov [3].

In this paper, we are concerned with *partitions* of H-free graphs such that *every* vertex class is either sparse or dense. It is easy to prove that H-free graphs can be partitioned into a bounded number of weakly ε-restricted subsets:

1.3 Theorem. *For every graph H, and all $\varepsilon > 0$, there is an integer N such that for every H-free graph G, there is a partition of $V(G)$ into at most N weakly ε-restricted subsets.*

This can be shown by applying 1.1 repeatedly to partition most of the vertices into weakly $\varepsilon/2$-restricted subsets, and then adding the remaining vertices into the largest set.

But what about partitions into sets that satisfy the stronger property of being ε-restricted? This is much harder, and the main result of this paper is the following:

1.4 Theorem. *For every graph H, and all $\varepsilon > 0$, there is an integer N such that for every H-free graph G, there is a partition of $V(G)$ into at most N ε-restricted subsets.*

This is significantly stronger than 1.3.

Here is a third statement midway between the last two: that under the same hypotheses, $V(G)$ is the union of at most a bounded number of ε-restricted subsets (not necessarily pairwise disjoint). This variation does not seem to be easy, although it does not imply 1.4 as far as we know.

Some remarks: sets of cardinality at most two are always ε-restricted, and for 1.4 it is sometimes necessary to use some ε-restricted subsets of cardinality at most two, even in graphs G with $|G|$

large. For example, let G be a star $K_{1,n}$ with n large, and let $\varepsilon < 1/3$: then every ε-restricted subset containing the centre of the star has cardinality at most two. (Note that this is not the case for 1.3; for example, a large star is already weakly ε-restricted.)

Second, our proof of 1.4 (and the proof in [6] of 2.2, which we will need to apply) does not use the regularity lemma. Thus, we anticipate that the number N in 1.4 is significantly smaller (as a function of $1/\varepsilon$) than numbers that are produced via the regularity lemma, but we have not made an estimate for it.

If $A, B \subseteq V(G)$ are disjoint, we say that B is ε-*sparse to* A (in G) if every vertex in B has at most $\varepsilon|A|$ neighbours in A; and B is ε-*dense to* A if B is ε-sparse to A in $\overline{G}$. The method of proof of 1.4 is via the following statement:

1.5 Theorem. *For every graph H, and all $\varepsilon, \eta, \theta > 0$, there exists an integer N such that, for every H-free graph G, there is a partition of $V(G)$ into nonempty sets*

$$A_1, \ldots, A_m, B_1, \ldots, B_m, C_1, \ldots, C_n,$$

where $m \leq |H|^2$ and $n \leq N$, such that:

- *$A_1, \ldots, A_m$ and $C_1, \ldots, C_n$ are ε-restricted sets;*

- *for $1 \leq i \leq m$, $|B_i| \leq \eta|A_i|$;*

- *for $1 \leq i \leq m$, B_i is either θ-sparse or θ-dense to A_i.*

We will prove this in section 2. In section 3 we prove another result, and combine these two to deduce 1.4.

Let us give an idea of how 1.5 will be used to prove 1.4. Let us say a "path-partition" is a sequence of $k+1$ disjoint subsets of $V(G)$, say $(W_0, \ldots, W_k)$, with three properties (that each of $W_0, \ldots, W_{k-1}$ is much bigger than W_k, that each of $W_0, \ldots, W_{k-1}$ is ε'-restricted for some appropriate ε', and that $W_{i+1} \cup \cdots \cup W_k$ is very sparse or very dense to W_i for each i). Note that the last term W_k need *not* be ε'-restricted; this is where the problem lies. We show in section 3 that if $(W_0, \ldots, W_k)$ is a path-partition with enough terms, then the union of all these sets can be partitioned into a small number of ε-restricted sets. This result is rather easy.

The role of 1.5 is to deduce something similar for successively shorter path-partitions. (If we can prove it for sequences of length one, then the main result follows, using the one-term sequence $V(G)$.) We will need to adjust the parameters of the sequence as its length k becomes smaller; that is, adjust the value of ε' such that $(W_0, \ldots, W_{k-1})$ are ε'-restricted, and adjust the density or sparsity condition.

Let $(W_0, \ldots, W_k)$ be the path-partition we want to handle now. Note that the induction means that we can handle all path-partitions that are *longer* than $(W_0, \ldots, W_k)$. We apply 1.5 to $G[W_k]$. That partitions W_k into a bounded number of ε-restricted sets and the pairs (A_s, B_s) $(1 \leq s \leq m)$ as in 1.5. We are happy with the ε-restricted sets, and we will use the pairs (A_s, B_s) to replace $(W_0, \ldots, W_k)$ by s longer path-partitions, disjoint and with union all the vertices in the union of $W_0, \ldots, W_k$ that are so far uncovered. To do so, we simply partition each W_j (for $0 \leq j < k$) into m large subsets W_i^s $(1 \leq s \leq m)$. Then we can apply the inductive hypothesis to the path-partition $\left(W_0^s, \ldots, W_{k-1}^s, A_s, B_s\right)$ for each s, and the result follows.

2 Proving the main lemma

In this section we prove 1.5. Let $A, B \subseteq V(G)$ be disjoint, and let $c, \varepsilon > 0$. We say that (A, B) is (c, ε)-*full* if for all $A' \subseteq A$ with $|A'| \geq c|A|$ and $B' \subseteq B$ with $|B'| \geq c|B|$, the number of edges between A', B' is at least $\varepsilon|A'| \cdot |B'|$. Similarly, (A, B) is (c, ε)-*empty* if it is (c, ε)-full in the complement graph. Thus, if (A, B) is (c, ε)-full, and $A' \subseteq A$ and $B' \subseteq B$ with $|A'|/|A|, |B'|/|B| \geq c' > c$ then (A', B') is $(c/c', \varepsilon)$-full.

We need a version of a standard result called the "embedding lemma":

2.1 Lemma. *Let G, H be graphs, let $0 < \varepsilon \leq 1/2$, and let A_v ($v \in V(H)$) be pairwise disjoint nonempty subsets of $V(G)$, such that for all distinct $u, v \in V(H)$, if u, v are adjacent in H then (A_u, A_v) is $(\varepsilon^{|H|}, \varepsilon)$-full, and if u, v are nonadjacent then (A_u, A_v) is $(\varepsilon^{|H|}, \varepsilon)$-empty. Then for each $v \in V(H)$ there exists $a_v \in A_v$ such that the map sending v to a_v for each $v \in V(H)$ is an isomorphism from H to an induced subgraph of G.*

Proof. We proceed by induction on $|H|$. If $|H| \leq 1$ the result is true, so we assume $|H| > 1$. Let $v \in V(H)$, and let N, M be the sets of neighbours of v in H and in $\overline{H}$ respectively. Let $c = \varepsilon^{|H|}$. For each $u \in N$ there are fewer than $c|A_v|$ vertices in A_v with fewer than $\varepsilon|A_u|$ neighbours in A_u, since (A_v, A_u) is (c, ε)-full; and similarly for each $u \in M$ there are fewer than $c|A_v|$ vertices in A_v with fewer than $\varepsilon|A_u|$ non-neighbours in A_u. Since $(|H| - 1)c < 1$ (because $\varepsilon \leq 1/2$), there exists $a_v \in A_v$ with at least $\varepsilon|A_u|$ neighbours in A_u for each $u \in N$, and at least $\varepsilon|A_u|$ non-neighbours in A_u for each $u \in M$. For each $u \in N$ let B_u be the set of neighbours of v in A_u, and for each $u \in M$ let B_u be the set of non-neighbours of v in A_u. Thus, each $B_u \neq \emptyset$, since $|B_u| \geq \varepsilon|A_u|$. Let H' be obtained from H by deleting v.

Thus, for all distinct $u, w \in V(H')$, if u, w are adjacent then (B_u, B_w) is $(c\varepsilon^{-1}, \varepsilon)$-full, and if u, w are nonadjacent then (B_u, B_w) is $(c\varepsilon^{-1}, \varepsilon)$-empty. From the inductive hypothesis, for each $u \in V(H')$ there exists $a_u \in B_u \subseteq A_u$ such that the map sending u to a_u for each $u \in V(H')$ is an isomorphism from H' to an induced subgraph of G. But then the theorem holds. This proves 2.1. ∎

The following is proved (without using the regularity lemma) in [6], theorem 2.2:

2.2 Lemma. *For all $c, \varepsilon, \tau > 0$ with $\varepsilon < \tau \leq 8/9$, there exists $\gamma > 0$ with the following property. Let G be a bipartite graph with a bipartition (A, B), with at least $\tau|A| \cdot |B|$ edges and with $A, B \neq \emptyset$. Then there exist $A' \subseteq A$ and $B' \subseteq B$ with $|A'|/|A|, |B'|/|B| \geq \gamma$, such that (A', B') is (c, ε)-full.*

Now we are ready to prove 1.5, but first let us sketch its proof. We want to partition $V(G)$ into a bounded number of ε-restricted sets and "θ-restricted pairs", by which we mean pairs of disjoint sets (A, B) where A is ε-restricted, B is much smaller than A, and B is either θ-dense or θ-sparse to A. Let $V(H) = \{v_1, \ldots, v_{|H|}\}$. Suppose that we have chosen pairwise disjoint, nonempty, subsets $D_1, \ldots, D_t$ of $V(G)$, such that:

- for all distinct i, j with $1 \leq i, j \leq t$, if v_i, v_j are adjacent in H then (D_i, D_j) is (x, y)-full, and if v_i, v_j are nonadjacent in H then (D_i, D_j) is (x, y)-empty, for some appropriate x, y.

Since G is H-free, we know that $t < |H|$, and our approach is to choose t as large as possible such that there is such a choice of $D_1, \ldots, D_t$. But to make use of the maximality of t, we also need that $D_1, \ldots, D_t$ are "not too small", and this is delicate. We cannot insist that they all have size some

3

constant times $|G|$, so let us see what we really need. To prove the theorem, we want a partition of the vertices not in $D_1 \cup \cdots \cup D_t$, using a bounded number of ε-restricted sets and θ-restricted pairs. Let us cover as much as we can, using (not too many) ε-restricted sets and θ-restricted pairs, and let E be the set of vertices that have not been covered. Now we can say what "not too small" means: we require that each of the sets $D_1, \ldots, D_t$ has size many times $|E|$. If $E \neq \emptyset$, we will show that we can choose another set D_{t+1} from within E, contrary to the maximality of t. Thus, $E = \emptyset$, and so we have the desired partition of $V(G)$.

Let us see this in more detail. We choose $t \leq |H|$ maximum such that there are vertex-disjoint subsets $D_1, \ldots, D_t$ of $V(G)$, and a bounded-size collection $\mathcal{A}$ of pairwise disjoint ε-restricted sets and θ-restricted pairs (the bound increasing with t), all disjoint from $D_1 \cup \cdots \cup D_t$, such that:

- For all distinct i, j with $1 \leq i, j \leq t$, if v_i, v_j are adjacent in H then (D_i, D_j) is (x, y)-full, and if v_i, v_j are nonadjacent in H then (D_i, D_j) is (x, y)-empty, for some appropriate x, y.

- $D_1, \ldots, D_t$ are nonempty and ε'-restricted where ε' is very small. (This extra condition is needed, because when we construct the new set D_{t+1}, parts of $D_1, \ldots, D_t$ will have to be discarded, and we need these parts to be ε-restricted so that we can add them to $\mathcal{A}$.)

- $D_1, \ldots, D_t$ are all many times bigger than E, where E is the "leftover" set E, that is, the set of vertices of G not in $D_1 \cup \cdots \cup D_t$ and not in any member of $\mathcal{A}$.

If $E = \emptyset$, we have proved what we want, so we suppose for a contradiction that $E \neq \emptyset$. We know that $t < |H|$, by 2.1, and now we will use E to try to increase t by 1. To build a new set D_{t+1} within E, we like vertices that are adjacent to at least a small (constant) fraction of the vertices in D_i for the values of i such that v_i, v_{t+1} are adjacent in H, and are nonadjacent to at least a small fraction of the vertices in D_i for the values of i such that v_i, v_{t+1} are nonadjacent (briefly, vertices that "have the desired adjacency"). But the vertices that do not have the desired adjacency are very sparse or very dense to some D_i, and so we can remove them all from E by adding a few more θ-restricted pairs to $\mathcal{A}$. (We have to keep all the sets and pairs of sets in $\mathcal{A}$ disjoint from $D_1, \ldots, D_t$, so we will have to shrink some of the sets a little, but that is straightfoward.) So we can assume that every vertex in E has the desired adjacency. By 1.2 we can choose a linear subset F_0 of E that is ε'-restricted, where ε' is very small. By applying 2.2 to each of the pairs F_0, D_i in turn, we can choose a linear subset D_{t+1} of F_0 that satisfies the first and second bullets above (changing ε' appropriately). We need to add a few sets to $\mathcal{A}$ to satisfy the third bullet. There are two main issues to worry about.

- First, when we applied 2.2 to F_0, D_i, the set D_i might shrink by a constant factor, and we need to take care of the "lost" vertices, those that belong to the old D_i and not the new one. But the old D_i was ε'-restricted where ε' is very small, and so we can arrange that the set of lost vertices is ε-restricted, by making sure it is not too small, and then we can add it to $\mathcal{A}$.

- Second, we have to arrange that the new leftover set, E' say, is small compared with $D_1, \ldots, D_{t+1}$. The sets $D_1, \ldots, D_t$ were shrunk in the process of finding D_{t+1}; but they remain at least a constant factor of their original sizes, and so their sizes are at least some constant times $|E|$. And the same is true for D_{t+1}, since D_{t+1} contains at least a linear fraction of F_0, and F_0 contains a linear fraction of E. But E' is a subset of E, so, while E' might be bigger than some of $D_1, \ldots, D_{t+1}$, its size is at most some large constant times the smallest of $D_1, \ldots, D_{t+1}$; and therefore, by repeatedly applying 1.2 and adding the sets we find to $\mathcal{A}$, we can reduce the size of E' by any constant factor that we wish, and so bring its size down to what we need.

We restate 1.5:

2.3 Theorem. *For every graph H, and all $0 < \varepsilon, \eta, \theta < 1$, there exists an integer N such that, for every H-free graph G, there is a partition of $V(G)$ into nonempty sets*

$$A_1, \ldots, A_m, B_1, \ldots, B_m, C_1, \ldots, C_n,$$

where $m \leq |H|^2$ and $n \leq N$, such that:

- *$A_1, \ldots, A_m$ and $C_1, \ldots, C_n$ are ε-restricted sets;*

- *for $1 \leq i \leq m$, $|B_i| \leq \eta|A_i|$; and*

- *for $1 \leq i \leq m$, B_i is either θ-sparse or θ-dense to A_i.*

Proof. We may assume that $\varepsilon, \eta, \theta < 1/3$, by reducing them if necessary. For each $\varepsilon' > 0$, let $\delta_{\varepsilon'}$ satisfy 1.2 with ε, δ replaced by $\varepsilon', \delta_{\varepsilon'}$.

Let $\varepsilon_{|H|} = \min(\varepsilon, (\theta/4)^{|H|})$. For $t = |H| - 1, |H| - 2, \ldots, 0$ in turn:

- for $i = t, t - 1, \ldots, 0$ in turn, let $\Gamma_{t,i} = \gamma_{t,i+1}\Gamma_{t,i+1}$ (or 1 if $i = t$); and choose $\gamma_{t,i}$ such that 2.2 holds, with $c, \varepsilon, \tau, \gamma > 0$ replaced by $\Gamma_{t,i}\varepsilon_{t+1}/3, \theta/4, \theta/2, \gamma_{t,i}$ respectively (by decreasing $\gamma_{t,i}$ if necessary we may assume that $\gamma_{t,i} \leq 1/3$ and $\gamma_{t,i} \leq \gamma_{t,i+1}$);

- let $\varepsilon_t = \gamma_{t,0}\varepsilon_{t+1}$.

For $1 \leq i \leq |H| + 1$, let $\varepsilon'_i = \varepsilon_i\Gamma_{i-1,0}$ and $\eta'_i = \min\left(\gamma_{i-1,0}, \frac{1}{2}\eta\delta_{\varepsilon'_i}\Gamma_{i-1,0}\right)$. For each $\gamma > 0$, let $\phi(\gamma)$ be the smallest nonnegative integer that satisfies $(1 - \delta_\varepsilon)^{\phi(\gamma)} \leq \gamma$. For $0 \leq t \leq |H|$, define

$$\ell_t = \sum_{1 \leq i \leq t} \phi(\eta'_i).$$

Let $N = \ell_{|H|} + |H|(|H| + 1)/2$; we claim that N satisfies the theorem.

Let G be H-free.

(1) *For all γ with $0 < \gamma < 1$, and for every $X \subseteq V(G)$, there is a partition of X into at most $\phi(\gamma) + 1$ sets, so that one of them has cardinality at most $\gamma|X|$ and the others are all ε-restricted.*

Let $X \subseteq V(G)$. Choose an ε-restricted set $A_1 \subseteq X$ with $|A_1| \geq \delta_\varepsilon|X|$; and inductively for each $i > 1$, choose an ε-restricted set $A_i \subseteq X \setminus (A_1 \cup \cdots \cup A_{i-1})$ with $|A_i| \geq \delta_\varepsilon|X \setminus (A_1 \cup \cdots \cup A_{i-1})|$. It follows that $|X \setminus (A_1 \cup \cdots \cup A_i)| \leq (1 - \delta_\varepsilon)^i|X|$ for each $i \geq 0$, and in particular when $i = \phi(\gamma)$. This proves (1).

Let $V(H)$ have vertices $v_1, \ldots, v_{|H|}$. For $0 \leq t \leq |H|$, we are interested in partitions of $V(G)$ into (possibly empty) sets $A_1, \ldots, A_m, B_1, \ldots, B_m, C_1, \ldots, C_\ell, D_1, \ldots, D_t$, and E, with the following properties:

- $m \leq t(t - 1)/2$ and $\ell \leq \ell_t$;

- $A_1, \ldots, A_m, C_1, \ldots, C_\ell$ and $D_1, \ldots, D_t$ are all nonempty;

- $A_1, \ldots, A_m, C_1, \ldots, C_\ell$ are ε-restricted;

- for $1 \le i \le m$, $|B_i| \le \eta |A_i|$, and B_i is either θ-sparse or θ-dense to A_i;

- $D_1, \ldots, D_t$ are ε_t-restricted;

- for $1 \le i < j \le t$, if v_i, v_j are adjacent in H then (D_i, D_j) is $(\varepsilon_t, \theta/4)$-full, and if v_i, v_j are nonadjacent then (D_i, D_j) is $(\varepsilon_t, \theta/4)$-empty;

- $|E| \le (\eta/2) \min(|D_1|, \ldots, |D_t|)$ if $t > 0$.

Let us call such a thing a *partition of type* (m, ℓ, t). To make clear which set plays which role in the partition, we will write them as:

$$(A_1, B_1), \ldots, (A_m, B_m)$$
$$C_1, \ldots, C_\ell$$
$$D_1, \ldots, D_t$$
$$E.$$

Choose such a partition, of type (m, ℓ, t) say, with $t \le |H|$ maximum. (This is possible, since G admits a partition of type $(0, 0, 0)$, setting $E = V(G)$.)

Since $\varepsilon_{|H|} \le (\theta/4)^{|H|}$, and $D_1, \ldots, D_t$ are nonempty, it follows from 2.1 that $t \le |H| - 1$. Choose pairwise disjoint subsets $E_1, \ldots, E_t$ of E with maximal union, such that for $1 \le i \le t$, if v_{t+1}, v_i are adjacent in H then E_i is $\theta/2$-sparse to D_i, and if v_{t+1}, v_i are nonadjacent in H then E_i is $\theta/2$-dense to D_i. Let $E_0 = E \setminus (E_1 \cup \cdots \cup E_t)$. Thus, for $1 \le i \le t$, E_0 is $(1 - \theta/2)$-dense to D_i if v_{t+1}, v_i are adjacent in H, and E_0 is $(1 - \theta/2)$-sparse to D_i if v_{t+1}, v_i are nonadjacent. Suppose, for a contradiction, that $E_0 \ne \emptyset$.

We recall that $|E| \le (\eta/2) \min(|D_1|, \ldots, |D_t|)$, and since $E \ne \emptyset$ (because $E_0 \subseteq E$), it follows that $|D_i| \ge \eta^{-1} > 1$ for $1 \le i \le t$. Thus, $\lfloor |D_i|/2 \rfloor \ge |D_i|/3$, for $1 \le i \le t$.

We recall that $\varepsilon'_{t+1} = \varepsilon_{t+1} \Gamma_{t,0}$; let $\delta'_{t+1} = \delta_{\varepsilon'_{t+1}}$. From 1.2 there is an ε'_{t+1}-restricted subset $F_0 \subseteq E_0$ with $|F_0| \ge \delta'_{t+1} |E_0|$. For $1 \le i \le t$ define $F_i \subseteq F_{i-1}$ with $|F_i| \ge \gamma_{t,i} |F_{i-1}|$, and $H_i \subseteq E_i$ with $|D_i|/2 \ge |H_i| \ge \gamma_{t,i} |D_i|$, as follows. Let us assume that v_{t+1}, v_i are adjacent (if they are non-adjacent, the construction is the same in the complement). Thus, F_{i-1} is $(1 - \theta/2)$-dense to D_i. (We remark that this is a weak assertion: it means that each vertex in F_{i-1} has at most $(1 - \theta/2)|D_i|$ non-neighbours in D_i, but θ may be very small.) From the definition of $\gamma_{t,i}$, there exist $F_i \subseteq F_{i-1}$ and $H'_i \subseteq D_i$, with $|F_i| \ge \gamma_{t,i} |F_{i-1}|$ and $|H'_i| \ge \gamma_{t,i} |D_i|$, such that (F_i, H'_i) is $(\Gamma_{t,i} \varepsilon_{t+1}/3, \theta/4)$-full. Let $H_i \subseteq H'_i$ of cardinality $\min(|H'_i|, \lfloor |D_i|/2 \rfloor)$. Thus, $|H_i| \ge \gamma_{t,i} |D_i|$, because either $|H_i| = |H'_i| \ge \gamma_{t,i} |D_i|$, or $|H_i| = \lfloor |D_i|/2 \rfloor \ge |D_i|/3 \ge \gamma_{t,i} |D_i|$. Since $|H_i| \ge |H'_i|/3$, it follows that (F_i, H_i) is $(\Gamma_{t,i} \varepsilon_{t+1}, \theta/4)$-full. This completes the inductive definition.

Thus, for $1 \le i \le t$, (F_i, H_i) is $(\Gamma_{t,i} \varepsilon_{t+1}, \theta/4)$-full if v_i, v_{t+1} are adjacent, and $(\Gamma_{t,i} \varepsilon_{t+1}, \theta/4)$-empty if v_i, v_{t+1} are non-adjacent. Also, since $|F_i| \ge \gamma_{t,i} |F_{i-1}|$ for $1 \le i \le t$, it follows that $|F_t| \ge \Gamma_{t,i} |F_i|$. Consequently (F_t, H_i) is $(\varepsilon_{t+1}, \theta/4)$-full if v_i, v_{t+1} are adjacent, and $(\varepsilon_{t+1}, \theta/4)$-empty if v_i, v_{t+1} are non-adjacent.

Now $|E| \le (\eta/2) \min(|D_1|, \ldots, |D_t|)$. We recall that $\eta'_{t+1} = \min\left(\gamma_{t,0}, \frac{1}{2} \eta \delta'_{t+1} \Gamma_{t,0}\right)$. By (1) there exist pairwise disjoint, nonempty, ε-restricted subsets $J_1, \ldots, J_n$ of $E_0 \setminus F_t$, with $n \le \phi(\eta'_{t+1})$, such that their union (J say) satisfies

$$|E_0 \setminus (F_t \cup J)| \le \eta'_{t+1} |E_0 \setminus F_t|.$$

We claim that the sets

$$(A_1, B_1), \ldots, (A_m, B_m), (D_1 \setminus H_1, E_1), \ldots, (D_t \setminus H_t, E_t)$$

$$C_1, \ldots, C_\ell, J_1, \ldots, J_n$$

$$H_1, \ldots, H_t, F_t$$

$$E_0 \setminus (F_t \cup J)$$

form a partition of $V(G)$ of type $(m + t, \ell + n, t + 1)$. To show this, we must check the following conditions, where $H_{t+1} = F_t$:

- Is it true that $m + t \leq t(t+1)/2$ and $\ell + n \leq \ell_{t+1}$? The first holds since $m \leq t(t-1)/2$; and the second holds since $\ell + n \leq \ell_t + \phi(\eta'_{t+1}) = \ell_{t+1}$.

- Is it true that $A_1, \ldots, A_m, D_1 \setminus H_1, \ldots, D_t \setminus H_t, C_1, \ldots, C_\ell, J_1, \ldots, J_n, H_1, \ldots, H_t$ and F_t are all nonempty? Certainly $A_1, \ldots, A_m, C_1, \ldots, C_\ell$ are nonempty from their definition, and so are $J_1, \ldots, J_n$. For $1 \leq i \leq t$, since $|E| \leq (\eta/2)|D_i|$ and $E \neq \emptyset$, it follows that $|D_i| \geq 2$; and so $D_i \setminus H_i \neq \emptyset$, since $|H_i| \leq |D_i|/2$. Also $|H_i| \geq \gamma_{t,i}|D_i| > 0$, so H_i is nonempty. Finally, $|F_t| \geq \Gamma_{t,0}|F_0|$, and $|F_0| \geq \delta'_{t+1}|E_0|$, and $E_0 \neq \emptyset$ by assumption; so $F_t \neq \emptyset$.

- Is it true that $A_1, \ldots, A_m, D_1 \setminus H_1, \ldots, D_t \setminus H_t, C_1, \ldots, C_\ell, J_1, \ldots, J_n$ are ε-restricted? $A_1, \ldots, A_m$, $C_1, \ldots, C_\ell$ and $J_1, \ldots, J_n$ are ε-restricted from their definition. For $1 \leq i \leq t$, D_i is ε_t-restricted, and since $|H_i| \leq |D_i|/2$, it follows that $D_i \setminus H_i$ is $2\varepsilon_t$-restricted and hence ε-restricted.

- Is it true that for $1 \leq i \leq m$, $|B_i| \leq \eta|A_i|$, and B_i is either θ-sparse or θ-dense to A_i; and for $1 \leq i \leq t$, $|E_i| \leq \eta|D_i \setminus H_i|$, and E_i is either θ-sparse or θ-dense to $D_i \setminus H_i$? The first is true from their definition. For the second, let $1 \leq i \leq t$. Then

$$|E_i| \leq |E| \leq (\eta/2)|D_i| \leq \eta|D_i \setminus H_i|$$

since $|D_i \setminus H_i| \geq |D_i|/2$. Also, E_i is either $\theta/2$-sparse to D_i (if v_i, v_{t+1} are adjacent in H) or $\theta/2$-dense to D_i (if v_i, v_{t+1} are nonadjacent); and so E_i is either θ-sparse or θ-dense to $D_i \setminus H_i$.

- Is it true that $H_1, \ldots, H_t, F_t$ are ε_{t+1}-restricted? For $1 \leq i \leq t$, D_i is ε_t-restricted, and since $|H_i| \geq \gamma_{t,i}|D_i| \geq \gamma_{t,0}|D_i|$, H_i is $\varepsilon_t/\gamma_{t,0}$-restricted and hence ε_{t+1}-restricted. Also, F_0 is ε'_{t+1}-restricted, and $|F_t| \geq \Gamma_{t,0}|F_0|$; and so F_t is $\varepsilon'_{t+1}/\Gamma_{t,0}$-restricted and hence ε_{t+1}-restricted.

- Is it true that for $1 \leq i < j \leq t+1$, if v_i, v_j are adjacent in H then (H_i, H_j) is $(\varepsilon_{t+1}, \theta/4)$-full, and if v_i, v_j are nonadjacent then (H_i, H_j) is $(\varepsilon_{t+1}, \theta/4)$-empty? If $j = t+1$, we already saw that (F_t, H_i) is $(\varepsilon_{t+1}, \theta/4)$-full if v_i, v_{t+1} are adjacent, and $(\varepsilon_{t+1}, \theta/4)$-empty if v_i, v_{t+1} are non-adjacent. So we may assume that $j \leq t$. Assume that v_i, v_j are adjacent (the other case is similar). Then (D_i, D_j) is $(\varepsilon_t, \theta/4)$-full, and since $|H_i| \geq \gamma_{t,0}|D_i|$ and $|H_j| \geq \gamma_{t,0}|D_j|$, and $\varepsilon_t = \gamma_{t,0}\varepsilon_{t+1}$, it follows that (H_i, H_j) is $(\varepsilon_{t+1}, \theta/4)$-full.

- Is it true that $|E_0 \setminus (F_t \cup J)| \leq (\eta/2)\min(|H_1|, \ldots, |H_t|, |F_t|)$? For $1 \leq i \leq t$, the choice of J implies that

$$|E_0 \setminus (F_t \cup J)| \leq \eta'_{t+1}|E_0 \setminus F_t|;$$

7

and
$$\eta'_{t+1}|E_0 \setminus F_t| \leq \gamma_{t,0}|E_0|,$$

since $\eta'_{t+1} \leq \gamma_{t,0}$. But
$$\gamma_{t,0}|E_0| \leq \gamma_{t,0}|E| \leq (\gamma_{t,0}\eta/2)|D_i| \leq (\eta/2)|H_i|,$$

since $|E| \leq (\eta/2)|D_i|$ and $\gamma_{t,0}|D_i| \leq |H_i|$. It follows that
$$|E_0 \setminus (F_t \cup J)| \leq (\eta/2)|H_i|$$

as claimed. Finally, to show that $|E_0 \setminus (F_t \cup J)| \leq (\eta/2)|F_t|$, observe that
$$|E_0 \setminus (F_t \cup J)| \leq \eta'_{t+1}|E_0| \leq \eta'_{t+1}|F_0|/\delta'_{t+1} \leq (\eta/2)\Gamma_{t,0}|F_0| \leq (\eta/2)|F_t|.$$

This proves that G admits a partition of type $(m+t, \ell+n, t+1)$, contrary to the choice of t, and so completes the proof that $E_0 = \emptyset$.

By renumbering, we may assume that $B_1, \ldots, B_r \neq \emptyset$, and $B_{r+1}, \ldots, B_m = \emptyset$, and $E_1, \ldots, E_s \neq \emptyset$, and $E_{s+1}, \ldots, E_t = \emptyset$. We claim that the pairs (A_i, B_i) for $1 \leq i \leq r$, the pairs (D_i, E_i) for $1 \leq i \leq s$, the sets A_i for $r+1 \leq i \leq m$, the sets D_i for $s+1 \leq i \leq t$, and the sets $C_1, \ldots, C_\ell$, satisfy the theorem. To show this, we observe:

- The sets
$$A_1, \ldots, A_m, B_1, \ldots, B_r, C_1, \ldots, C_\ell, D_1, \ldots, D_t, E_1, \ldots, E_s$$

  are pairwise disjoint and nonempty, and have union $V(G)$.

- The sets $A_1, \ldots, A_m, C_1, \ldots, C_\ell$ and $D_1, \ldots, D_t$ are ε-restricted (because each D_i is ε_t-restricted, and $\varepsilon_t \leq \varepsilon$).

- $|B_i| \leq \eta|A_i|$ for $1 \leq i \leq r$, and $|E_i| \leq |E| \leq (\eta/2)\min(|D_1|, \ldots, |D_t|) \leq \eta|D_i|$ for $1 \leq i \leq s$.

- For $1 \leq i \leq m$, B_i is either θ-sparse or θ-dense to A_i, and for $1 \leq i \leq s$, E_i is either $\theta/2$-sparse or $\theta/2$-dense to D_i, and hence either θ-sparse or θ-dense to D_i.

- $r + s \leq |H|^2$, since $r \leq m \leq t(t-1)/2$ and $s \leq t$, and $t \leq |H|$; and $\ell + (m-r) + (t-s) \leq N$, since
$$\ell \leq \ell_t \leq \ell_{|H|} = N - |H|(|H|+1)/2$$

  and
$$m - r + t - s \leq m + t \leq t(t-1)/2 + t \leq |H|(|H|+1)/2.$$

This proves 2.3. ∎

3 Path-partitions

We need the following two lemmas. For the first, see for example [1].

3.1 Lemma. *If $0 \leq k \leq n$ are integers, then $\binom{n}{k} \leq (en/k)^k$.*

The second lemma is the following (logarithms in this paper are to base e):

3.2 Lemma. *Let $\varepsilon > 0$ with $\varepsilon \leq 1/16$, and let $p \geq 0$ be an integer. Let G be a graph, and let A, B be nonempty disjoint subsets of $V(G)$, such that B is ε-sparse to A, and $\log(2|B|)/\varepsilon \leq p \leq |A|/12$. Then there exists $P \subseteq A$ with $|P| = p$, such that P is 2ε-sparse to B, and B is 12ε-sparse to P.*

Proof. We may assume that some vertex in B has a neighbour in A, because otherwise the result holds, and since $\varepsilon \leq 1/16$ it follows that $|A| \geq 16$. Let Q be the set of vertices in A with fewer than $2\varepsilon|B|$ neighbours in B, and let $q = |Q|$. There are at least $(|A| - q)(2\varepsilon|B|)$ and at most $\varepsilon|A| \cdot |B|$ edges between A and B, and so $q \geq |A|/2 \geq 8$. Let $k = \lceil 12\varepsilon p \rceil$.

Let $u_1, \ldots, u_{2p} \in Q$, not necessarily all distinct. Let y be the number of subsets of Q of cardinality p that contain all of $u_1, \ldots, u_{2p}$ (note that $p \leq |A|/12 \leq q$); and for each $v \in B$, let $z(v)$ be the number of subsets $I \subseteq \{1, \ldots, 2p\}$ of cardinality k such that u_i is adjacent to v for all $i \in I$ (note that $k = \lceil 12\varepsilon p \rceil \leq \lceil 2p \rceil = 2p$).

(1) *There is a choice of $u_1, \ldots, u_{2p}$ with $y = 0$ and $z(v) = 0$ for all $v \in B$.*

Choose $u_1, \ldots, u_{2p} \in Q$ uniformly and independently at random. Let $\overline{y}$ be the expectation of y, and $\overline{z}(v)$ the expectation of each $z(v)$. We will show that $\overline{y} < 1/2$, and $\overline{z}(v) \leq 1/(2|B|)$ for each $v \in B$, from which the claim follows. First,

$$\overline{y} = \binom{q}{p} \left(\frac{p}{q}\right)^{2p} \leq \left(\frac{ep}{q}\right)^p,$$

by 3.1. Since $p \leq |A|/12$ and $q \geq |A|/2$, it follows that $ep/q < 1/2$, and so $\overline{y} < 1/2$.

For $v \in B$, since v has at most $\varepsilon|A| \leq 2\varepsilon|Q|$ neighbours in Q, it follows that

$$\overline{z}(v) \leq \binom{2p}{k}(2\varepsilon)^k \leq \left(\frac{2ep}{k}\right)^k (2\varepsilon)^k = \left(\frac{4e\varepsilon p}{k}\right)^k \leq \left(\frac{e}{3}\right)^k \leq \left(\frac{e}{3}\right)^{12\varepsilon p}$$

from 3.1, and since $k \geq 12\varepsilon p$ and $e < 3$. From the hypothesis, $\log(2|B|) \leq \varepsilon p \leq 12\varepsilon p \log(3/e)$, and so $(e/3)^{12\varepsilon p} \leq 1/(2|B|)$. Hence $\overline{z}(v) \leq 1/(2|B|)$, and so the sum of $\overline{y}$ and all the $\overline{z}(v)$ ($v \in B$) is less than one. This proves (1).

Choose $u_1, \ldots, u_{2p}$ as in (1). Since $y = 0$ it follows that $|\{u_1, \ldots, u_{2p}\}| \geq p$; choose $P \subseteq \{u_1, \ldots, u_{2p}\}$ with $|P| = p$. Each vertex in P has at most $2\varepsilon|B|$ neighbours in B, since $P \subseteq Q$; and each $v \in B$ has at most $12\varepsilon p$ neighbours in P, since $z(v) = 0$. This proves 3.2. $\blacksquare$

Let G be a graph, let $k \geq 0$ be an integer, and let $\varepsilon > 0$. A (k, ε)-*path-partition* of G is a sequence $(W_0, W_1, \ldots, W_k)$ of subsets of $V(G)$, pairwise disjoint and with union $V(G)$, such that for $0 \leq i \leq k - 1$:

- W_i is ε-restricted;

- $|W_k| \leq |W_i|/12$;

- $W_{i+1} \cup \cdots \cup W_k$ is either $\varepsilon/12$-sparse or $\varepsilon/12$-dense to W_i.

If we are trying to partition $V(G)$ into ε-restricted sets, and G admits a (k, ε)-path-partition, then all but one of its sets are ε-restricted; the difficulty lies in handling the final set W_k.

3.3 Theorem. *Let $0 < \varepsilon \leq 1/3$, and let G be a graph admitting a $(k, \varepsilon/4)$-path-partition, where $k = \lceil 4/\varepsilon \rceil$. Then $V(G)$ can be partitioned into at most $2400\varepsilon^{-2}$ ε-restricted subsets.*

Proof. Let $(W_0, \ldots, W_k)$ be a $(k, \varepsilon/4)$-path-partition of G, let $p = |W_k|$, and $\varepsilon' = \varepsilon/48$.

(1) *We may assume that $\log(2kp) \leq \varepsilon'p$.*

Suppose not; then $\log(2kp) > \varepsilon p/48$, and since $k \leq 4/\varepsilon + 1 \leq 13/(3\varepsilon)$, it follows that

$$26p/(3\varepsilon) \geq 2kp > e^{\varepsilon p/48} \geq (\varepsilon p/48)^3/6,$$

(because $e^x \geq x^3/3!$ for all $x > 0$). We deduce that $p^2 \leq 52 \cdot 48^3/\varepsilon^4$, and so $p \leq 2398.5/\varepsilon^2$. Since $k \leq 13/(3\varepsilon) \leq 1.5/\varepsilon^2$, the theorem holds, because $V(G)$ is the union of $W_0, \ldots, W_{k-1}$ and the p singletons $\{v\}$ ($v \in W_k$). This proves (1).

(2) *For $0 \leq i \leq k$, there exists $C_i \subseteq W_i$ with $|C_i| = p$, such that for $0 \leq i \leq k - 1$, either*

- *C_i is $2\varepsilon'$-sparse to $C_{i+1} \cup \cdots \cup C_k$, and $C_{i+1} \cup \cdots \cup C_k$ is $12\varepsilon'$-sparse to C_i, or*

- *C_i is $2\varepsilon'$-dense to $C_{i+1} \cup \cdots \cup C_k$, and $C_{i+1} \cup \cdots \cup C_k$ is $12\varepsilon'$-dense to C_i.*

The choice of C_i is inductive, as follows: let $C_k = W_k$, and now suppose that $0 \leq i \leq k - 1$, and $C_{i+1}, \ldots, C_k$ are defined. Let $B = C_{i+1} \cup \cdots \cup C_k$. Thus, $|B| = (k - i)p$ and B is either ε'-sparse or ε'-dense to W_i (because $(W_0, \ldots, W_k)$ is a $(k, 12\varepsilon')$-path-partition). Moreover, $p = |W_k| \leq |W_i|/12$. Suppose first that B is ε'-sparse to W_i. By (1), $\log(2|B|) \leq \log(2kp) \leq \varepsilon'p$. By 3.2, taking $A = W_i$, and replacing ε by ε', we deduce that there exists $C_i \subseteq W_i$ with $|C_i| = p$, such that C_i is $2\varepsilon'$-sparse to B, and B is $12\varepsilon'$-sparse to C_i. Similarly, if B is ε'-dense to W_i, then 3.2 applied in $\overline{G}$ implies that there exists $C_i \subseteq W_i$ with $|C_i| = p$, such that C_i is $2\varepsilon'$-dense to B, and B is $12\varepsilon'$-dense to C_i. In either case, this completes the inductive definition of $C_0, \ldots, C_k$, and so proves (2).

Now for $0 \leq i \leq k - 1$, either C_i is $2\varepsilon'$-sparse to $C_{i+1} \cup \cdots \cup C_k$, or $2\varepsilon'$-dense to $C_{i+1} \cup \cdots \cup C_k$; choose $I \subseteq \{0, \ldots, k - 1\}$ with $|I| \geq k/2$ such that either C_i is $2\varepsilon'$-sparse to $C_{i+1} \cup \cdots \cup C_k$ for all $i \in I$, or C_i is $2\varepsilon'$-dense to $C_{i+1} \cup \cdots \cup C_k$ for all $i \in I$. Let $C = \bigcup_{i \in I \cup \{k\}} C_i$.

(3) *C is ε-restricted.*

To see this, suppose first that C_i is $2\varepsilon'$-sparse to $C_{i+1} \cup \cdots \cup C_k$ for all $i \in I$. Let $v \in C_j$ where $j \in I \cup \{k\}$, and let $I_1 = \{i \in I : i < j\}$, and $I_2 = \{i \in I \cup \{k\} : i > j\}$. Since C_j is $2\varepsilon'$-sparse to $C_{j+1} \cup \cdots \cup C_k$, it follows that v has at most $2\varepsilon'p(k - j) \leq \varepsilon p(k - j)/4$ neighbours in $C_{j+1} \cup \cdots \cup C_k$

10

(and hence at most the same number in $\bigcup_{i \in I_2} C_i$). For each $i \in I_1$, since $C_{i+1} \cup \cdots \cup C_k$ (and hence C_j) is $12\varepsilon'$-sparse to C_i, it follows that v has at most $12\varepsilon' p = \varepsilon p/4$ neighbours in C_i; and therefore v has at most $\varepsilon pj/4$ neighbours in $\bigcup_{i \in I_1} C_i$. Since v has at most p neighbours in C_j, it follows that v has at most

$$\varepsilon p(k-j)/4 + \varepsilon pj/4 + p = \varepsilon pk/4 + p \leq \varepsilon pk/2 \leq \varepsilon |C|$$

neighbours in C (here we use that $k \geq 4/\varepsilon$ and $|C| \geq pk/2$), and so C is ε-restricted. If C_i is $2\varepsilon'$-dense to $C_{i+1} \cup \cdots \cup C_k$ for all $i \in I$, we use the same argument in the complement. This proves (3).

For each $i \in I$, since $|C_i| = |W_k| \leq 3|W_i|/4$ and W_i is $\varepsilon/4$-restricted, it follows that $W_i \setminus C_i$ is ε-restricted. But then $V(G)$ admits a partition into the sets W_i ($i \in \{0, \ldots, k-1\} \setminus I$), the sets $W_i \setminus C_i$ ($i \in I$), and C, and these sets are all ε-restricted. This is a total of $k + 1 \leq 4/\varepsilon + 2$ sets, and $4/\varepsilon + 2 \leq 5/\varepsilon \leq 5/(3\varepsilon^2)$. This proves 3.3. ∎

Next we combine 1.5 and 3.3 to prove an analogue of 3.3 for shorter and shorter sequences, and hence to prove 1.4. We will show the following (the -1 at the end is for inductive purposes):

3.4 Theorem. *Let H be a graph, and let $h = |H|^2$. Let $0 < \varepsilon \leq 1/3$, and let $K = \lceil 4/\varepsilon \rceil$. Let N be as in 1.5, with $\varepsilon, \eta, \theta$ replaced by $\varepsilon/(4(2h)^K), 1/(3h), \varepsilon/(48(2h)^K)$ respectively. Let $0 \leq k \leq K$, and let G be an H-free graph admitting a $(k, (2h)^{k-K}\varepsilon/4)$-path-partition $(W_0, \ldots, W_k)$. Then $V(G)$ can be partitioned into at most $h^{K-k}(2400/\varepsilon^2 + N) - N$ ε-restricted subsets.*

Proof. We may assume that $|H| \geq 2$ and so $h \geq 4$. We proceed by induction on $K - k$. If $K - k = 0$ then the result follows from 3.3, so we assume that $k < K$, and the result holds for $k + 1$. By 1.5, there is a partition of W_k into nonempty sets

$$A_1, \ldots, A_m, B_1, \ldots, B_m, C_1, \ldots, C_n,$$

where $m \leq h$ and $n \leq N$, such that:

- $A_1, \ldots, A_m$ and $C_1, \ldots, C_n$ are $\varepsilon/(4(2h)^K)$-restricted sets;

- for $1 \leq s \leq m$, $|B_s| \leq |A_s|/(3h)$;

- for $1 \leq s \leq m$, B_s is either $\varepsilon/(48(2h)^K)$-sparse or $\varepsilon/(48(2h)^K)$-dense to A_s.

Let X be the union of the sets $C_1, \ldots, C_n$. Since each of these sets is $\varepsilon/(4(2h)^K)$-restricted and hence ε-restricted, it follows that X can be partitioned into at most N ε-restricted sets. If $m = 0$, the theorem holds, since $V(G)$ is the union of the sets $C_1, \ldots, C_n$ and $W_0, \ldots, W_{k-1}$, and $n \leq N$ and $k \leq 4/\varepsilon$; so we assume that $m > 0$. Consequently $|W_k| \geq |A_1| \geq 3h|B_1| \geq 3h$, and for $0 \leq i < k$, $|W_i| \geq 12|W_k| \geq 36h$; and hence $|W_i|/(2h) \geq 1$. It follows that $\lceil |W_i|/(2h) \rceil \leq |W_i|/h$, and therefore there are h pairwise disjoint subsets of W_i, each of cardinality at least $|W_i|/(2h)$. Consequently we may choose subsets $W_i^1, \ldots, W_i^m$ of W_i, pairwise disjoint and with union W_i, and each of cardinality at least $|W_i|/(2h)$.

(1) *For $1 \leq s \leq m$, $(W_0^s, \ldots, W_{k-1}^s, A_s, B_s)$ is a $(k+1, (2h)^{k+1-K}\varepsilon/4)$-path-partition of $G[V_s]$, where $V_s = W_0^s \cup \cdots \cup W_{k-1}^s \cup A_s \cup B_s$.*

To see this, we must show that

- A_s is $(2h)^{k+1-K}\varepsilon/4$-restricted;

- $|A_s| \geq 12|B_s|$;

- B_s is either $(2h)^{k+1-K}\varepsilon/48$-sparse or $(2h)^{k+1-K}\varepsilon/48$-dense to A_s;

and also that for $0 \leq i \leq k-1$:

- W_i^s is $(2h)^{k+1-K}\varepsilon/4$-restricted;

- $|W_i^s| \geq 12|B_s|$;

- $W_{i+1}^s \cup \cdots \cup W_{k-1}^s \cup A_s \cup B_s$ is either $(2h)^{k+1-K}\varepsilon/48$-sparse or $(2h)^{k+1-K}\varepsilon/48$-dense to W_i^s.

The first three statements are immediate from the definition of the pair (A_s, B_s). For the last three, let $0 \leq i \leq k-1$. It follows that W_i is $(2h)^{k-K}\varepsilon/4$-restricted, and since $|W_i^s| \geq |W_i|/(2h)$, we deduce that W_i^s is $(2h)^{k+1-K}\varepsilon/4$-restricted.

To show that $|W_i^s| \geq 12|B_s|$, observe that $|W_i| \geq 12|W_k| \geq 12|A_s| \geq 36h|B_s|$, and so $|W_i^s| \geq |W_i|/(2h) \geq 18|B_s|$.

Finally, to show that $W_{i+1}^s \cup \cdots \cup W_{k-1}^s \cup A_s \cup B_s$ is either $(2h)^{k+1-K}\varepsilon/48$-sparse or $(2h)^{k+1-K}\varepsilon/48$-dense to W_i^s, observe that, since $(W_0, \ldots, W_k)$ is a $\big(k, (2h)^{k-K}\varepsilon/4\big)$-path-partition, it follows that $W_{i+1} \cup \cdots \cup W_k$ is either $(2h)^{k-K}\varepsilon/48$-sparse or $(2h)^{k-K}\varepsilon/48$-dense to W_i, and hence so is

$$W_{i+1}^s \cup \cdots \cup W_{k-1}^s \cup A_s \cup B_s;$$

and therefore the latter is either $(2h)^{k+1-K}\varepsilon/48$-sparse or $(2h)^{k+1-K}\varepsilon/48$-dense to W_i^s, since $|W_i^s| \geq |W_i|/(2h)$. This proves (1).

From (1) and the inductive hypothesis, V_s can be partitioned into at most $h^{K-k-1}(2400/\varepsilon^2 + N) - N$ ε-restricted subsets, for $1 \leq s \leq m$. Since the sets $V_1, \ldots, V_m, C_1, \ldots, C_n$ are pairwise disjoint and have union $V(G)$, we deduce that $V(G)$ can be partitioned into at most

$$h(h^{K-k-1}(2400/\varepsilon^2 + N) - N) + N \leq h^{K-k}(2400/\varepsilon^2 + N) - N$$

ε-restricted subsets. This proves 3.4. ∎

To deduce 1.4, we may assume that $\varepsilon \leq 1/3$, by reducing ε if necessary; then 1.4 is immediate from 3.4 with $k = 0$, applied to the $\big(0, (2h)^{-K}\varepsilon/4\big)$-path-partition with one term $V(G)$.

Finally, there is a strengthening of 1.1 due to Nikiforov [5]:

3.5 Theorem. *For every graph H and all $\varepsilon > 0$, there exists $\delta > 0$ such that if G is a graph containing fewer than $(\delta|G|)^{|H|}$ induced copies of H, then there exists $S \subseteq V(G)$ with $|S| \geq \delta|G|$ such that $G[S]$ is weakly ε-restricted.*

As before, we can remove "weakly"; and this suggests that perhaps an analogue of 1.4 holds, with the "H-free" hypothesis replaced by the hypothesis of 3.5. This is false (take the union of a small random graph and a large stable set), but Tung Nguyen [4] has recently proved the following:

3.6 Theorem. *For every graph H, and all $\varepsilon > 0$, there exist $C > 0$ and an integer $N > 0$ such that for every graph G, if k denotes the number of distinct isomorphisms from H to induced subgraphs of G, then there exists $X \subseteq V(G)$ with $|X| \leq Ck^{1/|H|}$, and a partition of $V(G \setminus X)$ into at most N ε-restricted sets.*

His proof is by a modification of the arguments of this paper.

Acknowledgement

We are very grateful for a referee's report, which was immensely careful and helpful.

References

[1] B. Bollobás, *Modern Graph Theory*, page 216.

[2] N. Bousquet, A. Lagoutte and S. Thomassé, "The Erdős-Hajnal conjecture for paths and antipaths", *J. Combinatorial Theory, Ser. B* **113** (2015), 261–264.

[3] J. Fox and B. Sudakov, "Induced Ramsey-type theorems", *Advances in Math.* **219** (2008), 1771–1800.

[4] T. Nguyen, "A further extension of Rödl's theorem", submitted for publication, `arXiv:2208:07483`.

[5] V. Nikiforov, "Edge distribution of graphs with few copies of a given graph", *Combin. Probab. Comput.* **15** (2006), 895–902.

[6] A. Scott, P. Seymour and S. Spirkl, "Pure pairs. IV. Trees in bipartite graphs", *J. Combinatorial Theory, Ser. B* **161** (2023), 120–146, `arXiv:2009.09426`.

[7] V. Rödl, "On universality of graphs with uniformly distributed edges", *Discrete Math.* **59** (1986), 125–134.