Blow-up Lemmas for Sparse Graphs[*]

Peter Allen   Julia Böttcher   Hiệp Hàn [†]   Yoshiharu Kohayakawa [‡]

Yury Person [§]

Dedicated to Vojtěch Rödl, on the occasion of his 75th birthday.

Abstract: The blow-up lemma states that a system of super-regular pairs contains all bounded degree spanning graphs as subgraphs that embed into a corresponding system of complete pairs. This lemma has far-reaching applications in extremal combinatorics.

We prove sparse analogues of the blow-up lemma for subgraphs of random and of pseudorandom graphs. Our main results are the following three sparse versions of the blow-up lemma: one for embedding spanning graphs with maximum degree Δ in subgraphs of $G(n,p)$ with $p = C(\log n/n)^{1/\Delta}$; one for embedding spanning graphs with maximum degree Δ and degeneracy D in subgraphs of $G(n,p)$ with $p = C\left(\log n/n\right)^{1/(2D+1)}$; and one for embedding spanning graphs with maximum degree Δ in $(p, cp^{\max(4,(3\Delta+1)/2)}n)$-bijumbled graphs.

We also consider various applications of these lemmas.

Key words and phrases: regularity method, extremal problems in random and pseudorandom graphs

Contents

Chapter 1

Introduction, applications, results and proof overview

1.1 Introduction

Szemerédi's *regularity lemma* [77], originally developed for the proof of Szemerédi's celebrated result on arithmetic progressions [76], is one of the most influential tools in modern Discrete Mathematics. Numerous variants of this lemma, which is an approximate structure theorem for graphs, have been established for applications in other areas of mathematics, such as additive number theory, information theory, and statistical mechanics [21, 29].

Applications of the regularity lemma include a wealth of results in such diverse areas as extremal combinatorics, Ramsey theory, property testing, or discrete geometry. In such applications the regularity lemma is usually complemented by the *counting lemma* or the *blow-up lemma*. The former allows one to deduce estimates on small substructure counts from the (finitely sized) approximate structure provided by the regularity lemma. The latter, on the other hand, is powerful for establishing global structural properties. More precisely, the blow-up lemma, proved by Komlós, Sárközy and Szemerédi [56], permits the embedding of certain bounded degree spanning graphs. Alternative proofs can be found in [57, 71, 72]; for a nice introduction to the blow-up lemma and explanations about how it is used in applications see the surveys [55, 58].

One limitation of the original regularity lemma is that, because of the error terms, this lemma is suitable only for dense graphs, that is n-vertex graphs with $\Omega(n^2)$ edges. Nonetheless it is desirable to have equally effective tools at hand for sparse graphs. The most prominent example of why such sparse structures are of importance is without doubt the famous Green-Tao Theorem [43] on arithmetic progressions in the primes, which uses an approximate structure theorem for certain sparse subsets of the integers. Moreover, the modern branch of extremal combinatorics concerned with resilience results (a term coined by Sudakov and Vu [75]), which recently received much interest, investigates such sparse graphs.

Analogues of the regularity lemma that also work in a sparse setting, that is, for n-vertex graphs with $o(n^2)$ edges, do exist [50, 74]. However, a corresponding counting lemma simply fails to be true in general (see, e.g., [51]). This impediment can be overcome by posing additional restrictions on the graphs under

consideration. The existing counterexamples are known not to occur in random or certain pseudorandom graphs, and it was a major breakthrough when recently counting lemmas could finally be established in these settings: Counting lemmas for subgraphs of random graphs were proved in [20, 32, 73], and for subgraphs of pseudorandom graphs in [31].

What was so far missing in this effort to transfer these tools to the sparse setting was a sparse version of the blow-up lemma. An important step in this direction was taken in [54], where an embedding lemma for bounded degree graphs on cn vertices for some very small constant c in n-vertex subgraphs of random graphs was established. In [25] it was then shown that methods developed in [13] can be used to prove a blow-up type result for embedding almost spanning bipartite graphs. Moreover, in [19] a sparse embedding lemma for the special case of spanning triangle factors was proved. Analogues of these partial results for pseudorandom graphs are not known. But in [65] the importance of a generalisation of the blow-up lemma to pseudorandom graphs was acknowledged.

In this paper we provide this missing piece and establish several sparse versions of the blow-up lemma for random and for pseudorandom graphs. We also discuss a variety of relatively straightforward applications of these lemmas and indicate more intricate applications, which will appear elsewhere.

Organisation

We first motivate our blow-up lemmas in Section 1.2 by collecting various applications of these lemmas, some of which will be proved in Chapter 6 and some of which will be proved in future papers. In Section 1.3 we then provide our blow-up lemmas together with the necessary notation. We also state *regularity inheritance lemmas* which are necessary in applications. In Section 1.4 we provide an outline of the proofs of the blow-up lemmas. In Chapters 2–5 we give the proofs of our blow-up lemmas. We will describe the purposes of these various chapters in more detail in the proof outline (Section 1.4). We give the proofs of our applications in Chapter 6, and finish off with some concluding remarks in Chapter 7.

Notational remarks

We will routinely omit floor and ceiling signs when they do not affect the argument. All logarithms are taken to base 2.

1.2 Applications

In this section we collect a number of applications of our main results, establishing new structural properties of random and of pseudorandom graphs, and improving on a variety of earlier work in this area. We also present applications in Ramsey theory and the theory of positional games. We defer proofs of all these results to Chapter 6.

Before we provide the results, let us introduce the models of random and pseudorandom graphs that we use in this paper. The random graph model we work with is the binomial model $G(n, p)$, where each potential edge is included in a graph with n vertices independently with probability $p = p(n)$, and whose study was pioneered in an influential sequence of papers by Erdős and Rényi[1] (see, e.g., [24, 46] for the

---

[1]In fact Erdős and Rényi studied the related model $G(n,m)$.

background). If $G(n,p)$ has some property with probability tending to 1 as n tends to infinity, we say $G(n,p)$ has this property *asymptotically almost surely*, abbreviated a.a.s.

The study of pseudorandom graphs was initiated by Thomason [78], who asked for a set of easy deterministic properties enjoyed by $G(n,p)$ a.a.s. which by themselves imply many of the complex structural properties we know to hold for $G(n,p)$. The pseudorandomness notion we use is closely related to the notion suggested by Thomason, and is among the most widely used ones by now (for example, the sparse counting lemma in [31] is developed for this notion as well). We say a graph Γ is (p,β)-*bijumbled* if for all subsets $X, Y \subseteq V(\Gamma)$ we have

$$\left| e_\Gamma(X,Y) - p|X||Y| \right| \leq \beta \sqrt{|X||Y|}, \tag{1.1}$$

where $e_\Gamma(A,B)$ is the number of pairs in $A \times B$ which form edges in Γ. The random graph $G(n,p)$ is with high probability (p,β)-bijumbled with $\beta = O(\sqrt{pn})$, which justifies this definition.

Another class of pseudorandom graphs we shall refer to in the applications are (n,d,λ)-graphs. These have been studied extensively and are a special case of bijumbled graphs. For a graph Γ let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the eigenvalues of the adjacency matrix of Γ. We call $\lambda(\Gamma) := \max\{|\lambda_2|, |\lambda_n|\}$ the *second eigenvalue* of Γ. An (n,d,λ)-*graph* Γ is a d-regular graph on n vertices with $\lambda(\Gamma) \leq \lambda$. The connection between (n,d,λ)-graphs and bijumbled graphs is provided by the well-known *expander mixing lemma* (see, e.g., [18]), which states that if Γ is an (n,d,λ)-graph, then

$$\left| e_\Gamma(A,B) - \tfrac{d}{n}|A||B| \right| \leq \lambda(\Gamma)\sqrt{|A||B|}$$

for all disjoint subsets $A, B \subseteq V(\Gamma)$. This implies that Γ is $\left(\tfrac{d}{n}, \lambda(\Gamma)\right)$-bijumbled.

1.2.1 Universal graphs

It is interesting to ask when random, or quasirandom graphs are $\mathcal{H}$-universal for a family of graphs $\mathcal{H}$, and recently this class of questions enjoyed popularity for various families $\mathcal{H}$. We say G is $\mathcal{H}$-*universal* if $H \subseteq G$ for each $H \in \mathcal{H}$. Two particular classes of interest are $\mathcal{H}(n,\Delta)$, the n-vertex graphs with maximum degree Δ, and $\mathcal{H}(n,d,\Delta)$, the n-vertex graphs with maximum degree Δ and degeneracy d.

An easy corollary of our blow-up lemma for random graphs (Lemma 1.21) is that for $\Delta \geq 2$, the random graph $G(n,p)$ is a.a.s. $\mathcal{H}(n,\Delta)$-universal if $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$, reproving a result of Dellamonica, Kohayakawa, Rödl and Ruciński [35] (for $\Delta \geq 3$) and Kim and Lee [49] (for $\Delta = 2$). We omit this proof here; it follows along similar lines to that of Theorem 1.1 below. Very recently Ferber and Nenadov [39] were able to improve on these results, showing $p \geq \left(n^{-1}\log^3 n\right)^{\frac{1}{\Delta-1/2}}$ suffices for $\Delta \geq 3$.

Turning to d-degenerate graphs, one would expect that $G(n,p)$ is already $\mathcal{H}(n,d,\Delta)$-universal for some rather smaller p if d is much less than Δ. For $d = 1$, i.e. for bounded degree trees, Montgomery (see [68]) announced that one may take $p = C(\Delta)\log^2 n/n$, which is optimal up to log factors. For $d \geq 2$, the best previous result, due to Ferber, Nenadov and Peter [40] is that we can take $p = \omega\left(\Delta^{12} n^{-1/4d}\log^3 n\right)$. We prove the following strengthening[2], using our blow-up lemma for degenerate graphs (Lemma 1.23).

---

[2]In fact Ferber, Nenadov and Peter proved a universality result in terms of a constraint on the 'maximum average degree'; in the class of graphs with degeneracy d this quantity is between d and $2d$.

Theorem 1.1. *For each $d \geq 2$, $\Delta \in \mathbb{N}$ and each $\gamma > 0$ there exists C such that the random graph $G(n,p)$ is a.a.s.*

(a) $\mathcal{H}(n, d, \Delta)$*-universal if $p \geq C\left(\frac{\log n}{n}\right)^{1/(2d+1)}$, and*

(b) $\mathcal{H}((1-\gamma)n, d, \Delta)$*-universal if $p \geq C\left(\frac{\log n}{n}\right)^{1/(2d)}$.*

The almost spanning universality result with $p \geq C\left(\frac{\log n}{n}\right)^{1/2d}$ should be compared to the recent result of Conlon, Ferber, Nenadov and Škorić [30], who showed that $G(n,p)$ is $\mathcal{H}((1-\gamma)n, \Delta)$-universal when $p = \omega(n^{-1/(\Delta-1)} \log^5 n)$ for $\Delta \geq 3$. Our result is better when $\Delta \geq 2d+1$. Conlon and Nenadov (see [69, Theorem 3.6]) showed that the random graph $G(n,p)$ is universal for the class of d-degenerate graphs on $(1-\gamma)n$ vertices for $p \geq \left(\frac{C\log^2 n}{n \log\log n}\right)^{1/d}$ improving part *(b)* of our Theorem 1.1 to the essentially best possible bound on p.

Another well-studied question is how big $e(G)$ must be for a universal graph G. Alon and Capalbo showed that the correct answer is $\Theta\left(n^{2-2/\Delta}\right)$ in [12] for the the class $\mathcal{H}(n, \Delta)$, and that if one further insists that $v(G) = n$ then the correct answer is still $O\left(n^{2-2/\Delta} \log^{4/\Delta} n\right)$ in [11] (where the extra polylog-factor is believed to be unnecessary). Theorem 1.1 provides sparser $\mathcal{H}(n, d, \Delta)$-universal graphs when $d \leq \Delta/4$.

Finally, we are able to show that sufficiently pseudorandom graphs are $\mathcal{H}(n, \Delta)$-universal. This was first suggested as an important direction of investigation by Krivelevich, Sudakov and Szabó [65], who established a corresponding result for (n, d, λ)-graphs Γ in the special case that one wants to embed a spanning triangle factor in Γ. In earlier work [5] we were able to improve on their result, showing that there is $\varepsilon > 0$ such that $(p, \varepsilon p^{5/2} n)$-bijumbled n-vertex graphs G with minimum degree $\frac{1}{2}pn$ actually contain the square of a Hamilton cycle, which implies the theorem of Krivelevich, Sudakov and Szabó with a better bound on λ. Observe that some minimum degree condition is necessary in this result since bijumbled graphs may contain isolated vertices.

For the class $\mathcal{H}(n, \Delta)$, to the best of our knowledge the only existing universality result is that obtained by applying the original blow-up lemma of Komlós, Sárközy and Szemerédi [56], which is possible when $\beta \leq n^{-\Omega(\Delta^{-2})}$. Using our sparse blow-up lemma for pseudorandom graphs (Lemma 1.25) we can at least achieve the correct power of Δ in the exponent.

Theorem 1.2. *For each $\Delta \geq 2$ there exists $c > 0$ such that for any $p > 0$, if*

$$\beta \leq cp^{\max(4, 3\Delta/2+1/2)}n,$$

any n-vertex (p, β)-bijumbled graph G with $\delta(G) \geq \frac{1}{2}pn$ is $\mathcal{H}(n, \Delta)$-universal.

We note that Alon and Bourgain [10] observed that results of this type have applications in combinatorial number theory. More precisely, they showed that the Cayley sum-graphs G of any multiplicative subgroup U of a finite field $\mathbb{F}_q$, where $V(G) = U$ and $uv \in E(G)$ whenever $u + v \in U$, is an (n, d, λ)-graph with $n = |U|$ and $\lambda = \sqrt{q}$. Thus the above theorem shows that sufficiently large multiplicative subgroups of finite fields contain all bounded-degree additive patterns; see [5] or [10] for the corresponding definitions and a more detailed discussion.

1.2.2 Partition universality

It is also possible to apply our blow-up lemmas to derive certain Ramsey-type results. Given a set $\mathcal{H}$ of graphs, the graph G is *r-partition universal* for $\mathcal{H}$ if in any *r*-colouring of the edge-set $E(G)$ there is a colour class which is $\mathcal{H}$-universal. Kohayakawa, Rödl, Schacht and Szemerédi [54] showed that for each r and Δ there exists C such that if $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$ then $G(n,p)$ is a.a.s. *r*-partition universal for $\mathcal{H}(n,\Delta)$. This result is also an easy corollary of our blow-up lemma for random graphs. The proof is, again, omitted, and follows along the same lines as that of the following result. Moreover, using our blow-up lemma for degenerate graphs, we can provide better lower bounds for *r*-partition universality for $\mathcal{H}(n,d,\Delta)$ if d is small compared to Δ.

Theorem 1.3. *For each $r,d,\Delta \in \mathbb{N}$ there exists C so that if $p \geq C\left(\frac{\log n}{n}\right)^{1/(2d)}$, then a.a.s. $G(Cn,p)$ is r-partition universal for $\mathcal{H}(n,d,\Delta)$.*

We remark that a result of Conlon and Nenadov (see [69, Theorem 6.3]) improves the bound from [54] on p to $p \geq C\left(\frac{\log n}{n}\right)^{1/(\Delta-1/2)}$ for the class of triangle-free graphs of maximum degree Δ for $\Delta \geq 5$ and is somewhat complementary to our Theorem 1.3.[3]

Kohayakawa, Rödl, Schacht and Szemerédi [54] also asked if a sufficiently pseudorandom graph is *r*-partition universal for $\mathcal{H}(n,\Delta)$. Using our blow-up lemma for pseudorandom graphs (Lemma 1.25) we can answer this in the affirmative.

Theorem 1.4. *For each $r,\Delta \in \mathbb{N}$ there exists C such that if $p > 0$ and G is a graph on Cn vertices which is $\left(p, \frac{1}{C}p^{\max(4,(3\Delta+1)/2)}n\right)$-bijumbled, then G is r-partition universal for $\mathcal{H}(n,\Delta)$.*

Folkman [41] showed that for each k there are K_{k+1}-free graphs G which are *Ramsey* for K_k, that is, in any colouring of the edges of G with two colours there is a monochromatic copy of K_k. We note that using a proof similar to that of Theorem 1.3 we can show the following Folkman-type result (which though can also be proved using the embedding lemma from [54] instead of our blow-up lemma).

Theorem 1.5. *For each $r \in \mathbb{N}$ and $\Delta \geq 2$, there exists a graph which contains no subgraph isomorphic to $K_{2\Delta,2\Delta}$ and is r-partition universal for $\mathcal{H}(n,\Delta)$.*

We note that the identical method also proves the existence of a $K_{2\Delta}$-free graph which is *r*-partition universal for $\mathcal{H}(n,\Delta)$. However, one can more easily obtain graphs which are even $K_{\Delta+2}$-free and are *r*-partition universal for $\mathcal{H}(n,\Delta)$ by starting with a $K_{\Delta+2}$-free graph which is *r*-colour Ramsey for $K_{\Delta+1}$ and blowing it up by Cn for some large C. We would like to thank David Conlon for pointing this out to us.

1.2.3 Maker-Breaker games

The next application of our blow-up lemmas concerns positional games. Let us stress that the results we derive have, as such, nothing to do with random or pseudorandom graphs, and are in that sense surprising.

The *Maker-Breaker H-game on K_n with bias b* is the following game. Maker and Breaker take turns to colour the edges of K_n with respectively red and blue. In each turn, Maker colours one edge, while

---

[3]For newer developments, which appeared after the completion of our work, see [1, 33, 36].

Breaker colours b edges (and edges may not be recoloured). Maker's aim is to create a red copy of H, while Breaker's aim is to prevent Maker from creating a red H. This class of games has been studied extensively (see the book of Beck [22] and a recent survey of Krivelevich [61]). If Breaker wins the H-game on K_n with bias b, then obviously Breaker also wins with bias $b+1$; it follows that there is a *threshold bias b* for each game which is the smallest b such that Breaker wins the H-game.

Let us briefly explain how random graphs are connected to these games. If either player, playing a randomised strategy, wins a given game with positive probability against perfect play from the opponent, then they have a deterministic winning strategy (since this is a finite two-person game of perfect information without draws). Bednarska and Łuczak [23] used this observation to determine the order of magnitude of the threshold bias for the H-game on K_n for each fixed graph H. They showed that this bias is $\Theta(p^{-1})$, where p is the threshold probability such that $G(n,p)$ contains H robustly, in the sense that removing any $\varepsilon(H)pn^2$ edges fails to destroy all copies of H.

However, for sequences $H = (H_n)$ with vertex number depending on n, much less is known. The critical bias for the Hamiltonicity game, with $H = C_n$, was determined only recently by Krivelevich [60]. For more general bounded degree graphs it is only known that the critical bias tends to infinity as n tends to infinity [17][4]. Using a recent result of Ferber, Krivelevich and Naves [38], which informally shows that Maker can a.a.s. make a subgraph of $G(n,p)$ with minimum degree very close to pn, we can show that in fact Maker can win the $\mathcal{H}(n,\Delta)$-universality game on $K_{(1+\delta)n}$ (i.e., Maker's graph contains simultaneously each $H \in \mathcal{H}(n,\Delta)$) against a polynomially growing bias, and that this bias can be increased on the class $\mathcal{H}(n,d,\Delta)$ if d is small enough. For the subclasses $\mathcal{H}'(n,\Delta)$ and $\mathcal{H}'(n,d,\Delta)$ of triangle-free graphs in $\mathcal{H}(n,\Delta)$ and $\mathcal{H}(n,d,\Delta)$ respectively, we can even win on K_n.

Theorem 1.6. *For each $d,\Delta \in \mathbb{N}$ and $\delta > 0$ there exists $c > 0$ such that the following holds. With bias b Maker wins*

(*a*) *the $\mathcal{H}(n,\Delta)$-universality game on $K_{(1+\delta)n}$ if $b \leq c\left(\frac{n}{\log n}\right)^{1/\Delta}$,*

(*b*) *the $\mathcal{H}'(n,\Delta)$-universality game on K_n if $b \leq c\left(\frac{n}{\log n}\right)^{1/\Delta}$,*

(*c*) *the $\mathcal{H}(n,d,\Delta)$-universality game on $K_{(1+\delta)n}$ if $b \leq c\left(\frac{n}{\log n}\right)^{1/(2d)}$,*

(*d*) *the $\mathcal{H}'(n,d,\Delta)$-universality game on K_n if $b \leq c\left(\frac{n}{\log n}\right)^{1/(2d+1)}$.*

We note that the proof of Theorem 1.6 uses a randomised strategy for Maker which succeeds with high probability against any strategy of Breaker, but this strategy does not seem to be easy to derandomise. It would be interesting to find explicit deterministic Maker strategies for these games, even for substantially smaller bias.

In [6] we strengthen the above result, showing analogues of (*b*) and (*d*) for the $\mathcal{H}(n,\Delta)$-universality game and the $\mathcal{H}(n,d,\Delta)$-universality game, respectively. The proof also relies on our blow-up lemmas for random graphs, but is more involved.

Theorem 1.7 (Allen, Böttcher, Kohayakawa, Naves, Person [6])**.**
For each $d,\Delta \in \mathbb{N}$ and $\delta > 0$ there exists $c > 0$ such that the following holds. With bias b Maker wins

---

[4]In [17] it is only explicitly shown that the critical bias is at least two, but Krivelevich [59] has explained that the techniques there would work for any constant b if n is large enough.

(*a*) *the* $\mathcal{H}(n,\Delta)$-*universality game on* K_n *if* $b \leq c\left(\frac{n}{\log n}\right)^{1/\Delta}$,

(*b*) *the* $\mathcal{H}(n,d,\Delta)$-*universality game on* K_n *if* $b \leq c\left(\frac{n}{\log n}\right)^{1/(2d+1)}$.

1.2.4 Resilience for low-bandwidth graphs

Resilience results for random graphs concern the question of how robustly the random graph possesses certain properties. Sudakov and Vu [75] promoted the systematic study of problems of this type, which henceforth enjoyed particular popularity in extremal graph theory. Our blow-up lemmas are particularly useful for obtaining resilience results for spanning or almost spanning subgraphs of random graphs. In this section we provide an example.

The Bollobás-Komlós conjecture, now the Bandwidth Theorem, proved by Böttcher, Schacht and Taraz [28], generalises (up to a small error term) several embedding results for spanning graphs. To state it, we need the concept of *bandwidth* bw(H), defined as the smallest natural b such that there is a labelling of the vertices of H by $1,\ldots,v(H)$ such that $|i-j| \leq b$ whenever ij is an edge in that ordering.

Theorem 1.8 (Bandwidth Theorem [28]). *For each* $\gamma > 0$ *and* Δ *there exists* $\beta > 0$ *and* n_0 *such that for all* $n \geq n_0$ *the following holds. If* H *is an* n-*vertex graph with maximum degree* $\Delta(H) \leq \Delta$, *chromatic number* $\chi(H)$, *and bandwidth* bw(H) $\leq \beta n$, *and* G *is an* n-*vertex graph with minimum degree* $\delta(G) \geq \left(\frac{\chi(H)-1}{\chi(H)} + \gamma\right)n$, *then* H *is a subgraph of* G.

What makes this result valuable is that many important bounded degree graph classes have bandwidth $o(n)$, such as bounded degree planar graphs, or more generally any class of bounded degree graphs defined by some forbidden minor (see [27]).

Böttcher, Kohayakawa and Taraz [25] proved a sparse analogue of this theorem for almost-spanning bipartite graphs H, working in typical random graphs $G(n,p)$ with $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$. Huang, Lee and Sudakov [45] on the other hand proved an analogue which allows for non-bipartite H, but only works in random graphs with p constant. Using our blow-up lemma for random graphs we can prove the following common extension[5] of these results.

Theorem 1.9. *For each* $\gamma > 0$ *and* Δ *there exist* $\beta > 0$ *and* C *such that if* $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$, *then a.a.s.* $\Gamma = G(n,p)$ *has the following property. If* H *is a* $(1-\gamma)n$-*vertex graph with maximum degree* $\Delta(H) \leq \Delta$, *chromatic number* $\chi(H)$, *and bandwidth* bw(H) $\leq \beta n$, *and* G *is an* n-*vertex subgraph of* Γ *with minimum degree* $\delta(G) \geq \left(\frac{\chi(H)-1}{\chi(H)} + \gamma\right)pn$, *then* H *is a subgraph of* G.

In [3] the following strengthening of this result is proved, showing that we can actually have $v(H) = n - Cp^{-2}$, which is optimal, and for some H even $v(H) = n$.

Theorem 1.10 (Allen, Böttcher, Ehrenmüller and Taraz [3]). *For each* $\gamma > 0$ *and* Δ *there exist* $\beta > 0$ *and* C *such that if* $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$, *then a.a.s.* $\Gamma = G(n,p)$ *has the following property. If* H *is an* n-*vertex graph with maximum degree* $\Delta(H) \leq \Delta$, *chromatic number* $\chi(H)$, *bandwidth* bw(H) $\leq \beta n$, *such that at least* Cp^{-2} *vertices of* H *are not contained in a triangle, and* G *is an* n-*vertex subgraph of* Γ *with minimum degree* $\delta(G) \geq \left(\frac{\chi(H)-1}{\chi(H)} + \gamma\right)pn$, *then* H *is a subgraph of* G.

---

[5]Actually Huang, Lee and Sudakov allowed for spanning embeddings of some graphs H, so that the result here is not quite an extension of their theorem, though the result in [3] is.

Furthermore a similar resilience statement for sufficiently bijumbled graphs, and an improvement when H has small degeneracy, are obtained in [3]. The proofs of all these results rely again on our blow-up lemmas, but they are significantly harder to obtain than Theorem 1.9, whose proof is a good illustration of how one can apply our blow-up lemmas to obtain resilience results.

1.2.5 Robustness of the Bandwidth Theorem

Robustness is a measure of how strongly an extremal theorem holds (alternative to resilience) and was proposed by Krivelevich, Lee and Sudakov [63]. They showed that if G is a *Dirac graph*, that is, a graph with $\delta(G) \geq \frac{1}{2}v(G)$, then there is C such that for $p \geq C\frac{\log n}{n}$, a.a.s. the graph G_p obtained by keeping edges of G independently with probability p is Hamiltonian. Here we prove a similar extension of the Bandwidth Theorem.

Theorem 1.11. *For each $\gamma > 0$ and Δ there exist $\beta > 0$ and C such that if $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$, the following holds. If H is an n-vertex graph with maximum degree $\Delta(H) \leq \Delta$, chromatic number $\chi(H)$, and bandwidth $\mathrm{bw}(H) \leq \beta n$, and G is an n-vertex graph with minimum degree $\delta(G) \geq \left(\frac{\chi(H)-1}{\chi(H)} + \gamma\right)n$, then a.a.s. H is a subgraph of G_p.*

We give a fairly detailed sketch proof of this result in Section 6.5. We note that in [2] we propose a strengthening of Theorem 1.10 which implies that under the conditions of Theorem 1.11, G_p is in fact a.a.s. universal for all n-vertex graphs H with $\Delta(H) \leq \Delta$ and $\mathrm{bw}(H) \leq \beta n$.

1.3 Statements of the results

The full versions of our blow-up lemmas are technically complex and so, in order to make their statement more compact, we will introduce a number of lengthy definitions. To provide motivation for these definitions we first state a simplified version of our blow-up lemma for random graphs, which is useful only in a few applications.

We will then provide the full version of the blow-up lemma for $G(n, p)$ in Section 1.3.2, and turn to a blow-up lemma for embedding degenerate graphs into subgraphs of $G(n, p)$ in Section 1.3.3. In Section 1.3.4 we present the blow-up lemma for bijumbled graphs. Finally, in Section 1.3.5 we state two *regularity inheritance* lemmas which are often necessary in applications of our random graph blow-up lemmas.

1.3.1 A simplified version

The dense blow-up lemma is an embedding lemma for super-regular pairs. Before we can formulate it we need some definitions. Let $G = (V, E)$ be a graph and A and B be disjoint subsets of V. We also write $V(G)$ and $E(G)$ for the set of vertices and edges of G, respectively, and $v(G)$ and $e(G)$ for the sizes of these sets. For a vertex $v \in V$ we denote by $N_G(v; A)$ the set of *neighbours* of v in A, that is, those vertices $u \in A$ with $vu \in E$. We write $\deg_G(v; A)$ for its cardinality $|N_G(v; A)|$. We write $e_G(A)$ for the number of edges in G with both vertices in A and $e_G(A, B)$ for the number of edges in G with one vertex in A and one in B. Given $X \subseteq V(G)$, we write $N(X; A) := \bigcup_{v \in X} N_G(v; A)$ for the *joint neighbourhood* in A of the

vertices. We let $N_G^*(X;A) := \bigcap_{v \in X} N_G(v;A)$ denote the *common neighbourhood* in A of vertices from X, and write $\deg_G(X;A)$ for its size $\left|N_G^*(X;A)\right|$. We will often omit the set brackets in $N_G^*(\{v_1, \ldots, v_\ell\};A)$, writing instead $N_G^*(v_1, \ldots, v_\ell;A)$. If the graph G is clear from the context we sometimes omit it in the subscripts. Furthermore, if we omit the set A we intend $A = V(G)$.

The *density* of the pair (A,B) is $d_G(A,B) = e_G(A,B)/(|A||B|)$. Let $\varepsilon, d \geq 0$. The pair (A,B) is called (ε,d)-*regular* (in G) if we have $d(A',B') \geq d - \varepsilon$ for all $A' \subseteq A$ with $|A'| \geq \varepsilon|A|$ and $B' \subseteq B$ with $|B'| \geq \varepsilon|B|$.

Remark 1.12. Observe that this differs from the usual definition of ε-regularity in that we only require a lower bound on $d(A',B')$. This is sometimes called *dense* in the literature, in particular in [54]. Note also that with our definition the density of an (ε,d)-regular pair (A,B) is only lower-bounded by $d - \varepsilon$. This, though non-standard, has the advantage (over the usual definition where we would require the density of (A,B) to be lower-bounded by d) that large subpairs of (ε,d)-regular pairs are (ε',d)-regular (rather than $(\varepsilon',d-\varepsilon')$-regular). We have to use a different regularity parameter, which is unavoidable, but at least we can stick to d for density throughout, and we advocate the use of this definition in future.

We shall use (sparse versions of) this regularity concept whenever we work with random graphs in this paper. But for pseudorandom graphs we unfortunately need the usual regularity definition with an upper bound on $d(A',B')$ as well. Whenever the distinction between the two different concepts is essential, which is usually not the case, we will explicitly state which version we are using, calling the former 'lower-regularity' and the latter 'full-regularity'. In most of the paper however, we will refer to both versions as 'regularity'. This will help us to unify much of the proofs of the blow-up lemmas for random graphs and for pseudorandom graphs. Why it is necessary to use these two different regularity variants is explained in Section 1.3.4.

An (ε,d)-regular pair (A,B) is called (ε,d)-*super-regular* if for every $u \in A$ we have $\deg_G(u;B) \geq (d-\varepsilon)|B|$ and for every $v \in B$ we have $\deg_G(v;A) \geq (d-\varepsilon)|A|$. The dense blow-up lemma, first proved by Komlós, Sárközy and Szemerédi [56], then states the following. Let G be a graph formed from a collection of (ε,d)-super-regular pairs with density $d \gg \varepsilon$, and let G^* be obtained from G by replacing the super-regular pairs with complete bipartite graphs. If H is a graph with maximum degree Δ which embeds into G^*, then H embeds into G.

This notion of regularity is not meaningful for sparse graphs, because the definition above implies that *any* pair with $o(n^2)$ edges is regular (with $d = 0$). For obtaining a meaningful sparse version of regularity and super-regularity we need to relate the density of a pair to the overall density of the graph under study. This can be obtained by replacing the density in the definitions above with the p-density for a suitable p.

For $0 < p < 1$, we define the *p-density* of the pair (A,B) to be $d_{G,p}(A,B) = e_G(A,B)/(p|A||B|)$. In other words, we simply scale the usual density by p. Sparse regular pairs are then defined as follows. The pair (A,B) is (ε,d,p)-*regular* (in G) if we have $d_{G,p}(A',B') \geq d - \varepsilon$ for all $A' \subseteq A$ with $|A'| \geq \varepsilon|A|$ and $B' \subseteq B$ with $|B'| \geq \varepsilon|B|$. Similarly, we define a pair to be a sparse super-regular pair if it is a sparse regular pair and satisfies a minimum degree condition.

The setting we will work with in this subsection is when the graph G into which we want to embed is a subgraph of a random graph Γ. As mentioned above, for pseudorandom graphs we will work with a slightly different notion of regularity (see Section 1.3.4). The parameter p will usually be (a constant factor away from) the density of the ambient graph Γ.

Definition 1.13 (Sparse super-regularity). A pair (A,B) in $G \subseteq \Gamma$ is called (ε, d, p)-*super-regular* (in G) if it is (ε, d, p)-regular and for every $u \in A$ and $v \in B$ we have

$$\deg_G(u;B) > (d - \varepsilon) \max\{p|B|, \deg_\Gamma(u;B)/2\}. \tag{1.2}$$
$$\deg_G(v;A) > (d - \varepsilon) \max\{p|A|, \deg_\Gamma(v;A)/2\}.$$

We remark that the term $(d - \varepsilon)p|B|$ is a natural lower bound in the minimum degree condition (1.2) by the following fact, which easily follows from the definition of regularity (and which we shall use throughout the paper).

Fact 1.14. *Let (A,B) be an (ε, d, p)-regular pair. Then less than $\varepsilon|A|$ vertices of A have less than $(d - \varepsilon)p|B|$ neighbours in B.*

Observe though that in (1.2) we have the additional term $(d - \varepsilon)\deg_\Gamma(u;B)/2$, which dominates when u has an exceptionally high Γ-degree into B. The reason why this is necessary will become clear in a little while after we discuss the need for regularity inheritance. Note also that it follows from the definition above that in a super-regular pair (A,B) we have $\deg_\Gamma(u;B) \geq (d - \varepsilon)p|B|$ for each $u \in A$, since $G \subseteq \Gamma$.

Unfortunately, this straightforward extension of the concept of super-regular pairs to the sparse setting is not on its own enough for a sparse blow-up lemma. To see this, suppose A, B and C are equally-sized vertex sets, and each pair is super-regular. We would like to find that there is a spanning triangle factor, but it is possible that for some $a \in A$ there are no edges between $N_G(a;B)$ and $N_G(a;C)$. Similarly, if A, B, C, D form a 4-cycle of super-regular pairs in G, it is nevertheless possible that for some $a \in A$ there is no 4-cycle in G using a and one vertex of each of the other three sets. These problems do not occur for dense (ε, d)-super-regular pairs for the following reason. In the first example, since $N_G(a;B)$ is of size at least $(d - \varepsilon)|B|$ and similarly for $N_G(a;C)$, the so-called slicing lemma guarantees that $\big(N_G(a;B), N_G(a;C)\big)$ is $(2\varepsilon/d, d)$-regular, that is, it *inherits* regularity from the pair (B,C) and thus in particular contains edges. A similar argument deals with the second example. The slicing lemma can be generalised to the sparse setting, and again easily follows from the definition of regularity.

Lemma 1.15 (Slicing lemma). *Let (A,B) be an (ε, d, p)-regular pair and $A' \subseteq A$, $B' \subseteq B$ be sets of sizes $|A'| \geq \alpha|A|$, $|B'| \geq \alpha|B|$. Then (A', B') is $(\varepsilon/\alpha, d, p)$-regular.*

But this unfortunately does not solve the problem indicated above because the size of $N_G(a;B)$ is only lower bounded by $(d - \varepsilon)p|B|$, which can be tiny compared to B.

To remedy this, in our blow-up lemmas we will explicitly require that the pair $\big(N_\Gamma(a;B), C\big)$, and similarly $\big(N_\Gamma(a;B), N_\Gamma(a;C)\big)$, 'inherit' regularity, that is, they are themselves regular pairs. Since $\deg_G(a;B)$ is a large fraction of $\deg_\Gamma(a,B)$ by super-regularity, the slicing lemma does show that regularity of $\big(N_G(a;B), C\big)$ follows from regularity of $\big(N_\Gamma(a;B), C\big)$, and similarly for $\big(N_G(a;B), N_G(a;C)\big)$. Note that this is the point where we require the second term in (1.2).

Requiring this 'inheritance of regularity' is reasonable because it is known that in (ε, d, p)-regular pairs contained in random graphs Γ, for almost all vertices a, the pairs $\big(N_\Gamma(a;B), C\big)$ and similarly $\big(N_\Gamma(a;B), N_\Gamma(a;C)\big)$ do inherit sparse regularity in this way. This phenomenon was studied in [42, 52, 54]. In this paper (see Section 1.3.5), we prove this statement with tight bounds on what 'almost all vertices' means.

Definition 1.16 (Regularity inheritance). Let A, B and C be vertex sets in $G \subseteq \Gamma$, where A and B are disjoint and B and C are disjoint, but we do allow $A = C$. We say that (A, B, C) has *one-sided* (ε, d, p)-*inheritance* if for each $u \in A$ the pair $\big(N_\Gamma(u; B), C\big)$ is (ε, d, p)-regular.

If in addition A and C are disjoint, then we say that (A, B, C) has *two-sided* (ε, d, p)-*inheritance* if for each $u \in A$ the pair $\big(N_\Gamma(u; B), N_\Gamma(u; C)\big)$ is (ε, d, p)-regular.

In this definition we take neighbourhoods in the ambient graph Γ instead of G because this turns out to be easier to handle in applications (where we need to take care of the vertices in regular pairs that do not satisfy these inheritance conditions).

The simplified version of our sparse blow-up lemma in random graphs now states that if we are given an equitable vertex partition $V_1 \,\dot\cup\, \cdots \,\dot\cup\, V_r$ of a graph $G \subseteq \Gamma$ such that all pairs are super-regular and all triples have one- and two-sided regularity inheritance, then any r-partite graph H with bounded maximum degree that can be embedded into the complete r-partite graph with partition classes $V_1, \ldots, V_r$ can also be embedded into G. Here, a partition of a set is called *equitable* if each pair of partition classes differ in size by at most one.

Lemma 1.17 (Simple blow-up lemma for $G(n, p)$). *For all $\Delta \geq 2, r \in \mathbb{N}$ and $d > 0$ there exist $\varepsilon > 0$ and C such that if $p \geq C(\log n / n)^{1/\Delta}$, then the random graph $\Gamma = G(n, p)$ a.a.s. has the following property. Let $G \subseteq \Gamma$ and H with $\Delta(H) \leq \Delta$ be graphs on n vertices. Suppose that H has an r-colouring with colour classes $X_1, \ldots, X_r$ and that $V(G) = V_1 \,\dot\cup\, \cdots \,\dot\cup\, V_r$ is an equitable partition with $|V_i| = |X_i|$ for all $i \in [r]$ and such that the following conditions hold.*

   (i) (V_i, V_j) is (ε, d, p)-super-regular in G for each $i, j \in [r]$ with $i \neq j$.

   (ii) (V_i, V_j, V_k) has one-sided (ε, d, p)-inheritance for each $i, j, k \in [r]$ with $i \neq j$ and $j \neq k$.

   (iii) (V_i, V_j, V_k) has two-sided (ε, d, p)-inheritance for each $i, j, k \in [r]$ with $i \neq j$, $j \neq k$, and $k \neq i$.

Then H is a subgraph of G.

Let us briefly comment on the lower bound $p \geq C(\log n / n)^{1/\Delta}$ on the probability for which our result works. We do not believe that this bound is in general best possible, though for $\Delta \in \{2, 3\}$ it is optimal up to the log-factor. However, it matches the best known current lower bound [35] for p such that $G(n, p)$ is universal for bounded degree spanning H. This universality result is easily implied by Lemma 1.17 with $G = \Gamma$. The problem of improving on [35] has been prominent in random graph theory for a few years, and seems to be hard. For a more detailed discussion, see Section 7.2.1.

The restriction $\Delta \geq 2$ is necessary for our proof as written. The statement above is true for $\Delta = 1$, when in fact we do not require conditions (ii) or (iii), but to see this it is easiest to verify Hall's condition in G directly rather than to modify our proof.

What is the difference between this simplified blow-up lemma and our full-strength blow-up lemma for random graphs, Lemma 1.21 (to be introduced below)? Firstly, in the latter we do not require the partition of G to be equitable but allow the partition classes to differ in size by a constant factor.

Secondly, we do not require all pairs in the partition to be super-regular (or have regularity inheritance). Instead we will introduce the concept of a 'reduced graph' which encodes where we have super-regular pairs in our partition. In fact we will even have two reduced graphs R and $R' \subseteq R$ where the former represents regular pairs and the latter super-regular pairs, the reason for which will become clear later.

Thirdly, we do not require two-sided regularity inheritance everywhere in the partition of G, but only in certain cases where triangles of H need to be embedded. This is helpful in some applications, for

example in Theorem 1.6 we can use the Ferber-Krivelevich-Naves strategy [38] to win Maker-Breaker games with spanning bounded degree triangle-free graphs. This strategy does not allow Maker to win (for example) the spanning triangle factor game, ultimately because of a failure of two-sided regularity inheritance.

Fourthly, the ε we can choose for our full-scale blow-up lemma does not depend on r, but only on the maximum degree of the reduced graph R'. This is a difference also in the dense setting (that is, when $p = 1$) to the blow-up lemma of [56]. In typical applications of the latter it is necessary to apply this blow-up lemma several times to embed a spanning H since this lemma only applies to small subgraphs of the reduced graph (because ε depends on r). It is then necessary to use so-called image restrictions and some manual embedding to connect up the pieces of H, which technically complicates the proofs. Our blow-up lemma on the other hand avoids this and is formulated with the intention that in applications it typically only needs to be applied once, which should make it simpler to use.

Finally, our blow-up lemma permits so-called image restrictions. Roughly speaking, image restriction means specifying, for certain vertices of H, small subsets of G into which these vertices are to be embedded. As explained in the last paragraph we believe that in the dense setting one usually does not need these image restrictions with our blow-up lemma. However, in the sparse setting they can be useful, as there may be a few vertices which do not satisfy the super-regularity or inheritance conditions required by the blow-up lemma in any sparse-regular partition of G. Hence we need to embed some H-vertices on these vertices 'by hand' before applying the blow-up lemma. These pre-embedded H-vertices then create image restrictions for their neighbours. Indeed in [3] precisely this approach is used to prove Theorem 1.10.

1.3.2   Random graphs

As explained in the previous subsection one of the differences of our blow-up lemma to the simplified version stated there is that it uses reduced graphs to specify where edges can be embedded in the graph G (which is a subgraph of a random graph Γ). We will now first define this concept and explain what we require of the graph H that we want to embed to be 'compatible' with such reduced graphs.

The setting in which we work is as follows. Let G and H be two graphs, on the same number of vertices, given with partitions $\mathcal{V} = \{V_i\}_{i \in [r]}$ and $\mathcal{X} = \{X_i\}_{i \in [r]}$ of their respective vertex sets. We call the parts V_i of G *clusters*. We say that $\mathcal{V}$ and $\mathcal{X}$ are *size-compatible* if $|V_i| = |X_i|$ for all $i \in [r]$. Moreover, for $\kappa \geq 1$ we say that $(G, \mathcal{V})$ is *κ-balanced* if there exists $m \in \mathbb{N}$ such that we have $m \leq |V_i| \leq \kappa m$ for all $i, j \in [r]$ (and thus for all $X_i \in \mathcal{X}$ if $\mathcal{V}$ and $\mathcal{X}$ are size-compatible). Our goal will be to embed H into G respecting these partitions.

As mentioned before, we will have two reduced graphs R and $R' \subseteq R$, where R' represents super-regular pairs and R regular pairs. More precisely, we require the following properties of R and R' and the partitions $\mathcal{V}$ and $\mathcal{X}$ of G and H.

Definition 1.18 (Reduced graphs and one-sided inheritance)**.** Let R and R' be graphs on r vertices.
- $(H, \mathcal{X})$ is an *R-partition* if each part of $\mathcal{X}$ is nonempty, and whenever there are edges of H between X_i and X_j, the pair ij is an edge of R,
- $(G, \mathcal{V})$ is an *(ε, d, p)-regular R-partition* if for each edge $ij \in E(R)$ the pair (V_i, V_j) is (ε, d, p)-regular.

In this case we also say that R is a *reduced graph* of the partition $\mathcal{V}$.

- $(G, \mathcal{V})$ is (ε, d, p)-*super-regular on* R' if for every $ij \in E(R')$ the pair (V_i, V_j) is (ε, d, p)-super-regular.

Suppose now that $(G, \mathcal{V})$ is an (ε, d, p)-regular R'-partition.

- $(G, \mathcal{V})$ has *one-sided inheritance* on R' if (V_i, V_j, V_k) has one-sided (ε, d, p)-inheritance for every $ij, jk \in E(R')$.

We occasionally also use these concepts when we only work on an induced subgraph of G, that is, for a pair $(G, \mathcal{V}')$ where $\mathcal{V}'$ is a partition of a subset of the vertices of G.

As in Lemma 1.17, we do require the case $i = k$ in the definition of one-sided inheritance.

It remains to describe where in our partitions we require two-sided inheritance. For this we first need to define so-called 'buffer sets' of vertices in H, containing 'potential buffer vertices'. The purpose of these buffer sets is that a subset of these vertices, later to be called 'buffer vertices', will be embedded last (for more detailed explanations see Section 1.4). For this to work we require the edges emanating from these vertices and their neighbours to be assigned to the super-regular pairs given by R' (and not to other pairs of R). It is only when a potential buffer vertex is contained in a triangle that we require two-sided inheritance on R'.

The buffer sets can be chosen by the user of the blow-up lemma. Moreover, we stress that they do not make our blow-up lemma less powerful in the dense setting than the blow-up lemma of [56], because the latter requires all edges of H to be assigned to super-regular pairs.

Definition 1.19 (Buffer sets and two-sided inheritance). Suppose $R' \subseteq R$ are graphs on r vertices, and $(H, \mathcal{X})$ is an R-partition and $(G, \mathcal{V})$ a size-compatible (ε, d, p)-regular R-partition. We say the family $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ of subsets $\tilde{X}_i \subseteq X_i$ is an (α, R')-*buffer* for $(H, \mathcal{X})$ if

- $\tilde{X}_i \subseteq X_i$ and $|\tilde{X}_i| \geq \alpha |X_i|$ for all $i \in [r]$ and
- for each $i \in [r]$ and each $x \in \tilde{X}_i$, the first and second neighbourhood of x *go along* R', that is, for each $xy, yz \in E(H)$ with $y \in X_j$ and $z \in X_k$ we have $ij \in E(R')$ and $jk \in E(R')$.

We also call the vertices in $\tilde{\mathcal{X}}$ *potential buffer vertices*. Moreover, $(G, \mathcal{V})$ has *two-sided inheritance* on R' for $\tilde{\mathcal{X}}$ if

- (V_i, V_j, V_k) has two-sided (ε, d, p)-inheritance whenever there is a triangle $x_i x_j x_k$ in H with $x_i \in \tilde{X}_i$, $x_j \in X_j$, and $x_k \in X_k$.

We remark that we shall later also occasionally refer to the set of actual buffer vertices as buffer, when it is clear from the context which set we mean.

Finally, our blow-up lemma allows image restrictions. These generalise the image restrictions permitted in the dense blow-up lemma. However, in the sparse setting the necessary conditions become somewhat more involved. The idea is as follows. Suppose we wish to embed a graph H^* into a graph G^*. Unfortunately G^* does not meet the conditions of our blow-up lemma, typically because regularity inheritance fails. We find a subgraph G of G^* which does meet the conditions of our blow-up lemma, and 'pre-embed' some vertices of H^* onto the vertices $V(G^*) \setminus V(G)$. This leaves the induced subgraph H of H^* to embed into G. The image restrictions then originate from these pre-embedded vertices: If $x \in X_i \subseteq V(H)$ has neighbours $\{z_1, \ldots, z_\ell\}$ in $V(H^*) \setminus V(H)$ which are pre-embedded to $\{u_1, \ldots, u_\ell\} = J_x \subseteq V(\Gamma) \setminus V(G)$, then J_x restricts the embedding of x to $I_x = N^*_{G^*}(J_x; V_i)$. In the following definition we do not explicitly refer to the graphs H^* and G^*, but only to abstract restricting sets J_x, so that we do not need to include the graphs H^* and G^* in our blow-up lemma. For the same reason we take

neighbourhoods in Γ instead of G^* in this definition. In addition, to simplify notation, we define an image restriction set I_x for each vertex x of H. For most vertices x, however, this set is the trivial set $I_x = X_i$ where X_i is the part of $\mathcal{X}$ containing x.

Definition 1.20 (Image restrictions). Let R be a graph on r vertices, and $(H, \mathcal{X})$ be an R-partition and $(G, \mathcal{V})$ a size-compatible (ε, d, p)-regular R-partition, where $G \subseteq \Gamma$. Let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a collection of subsets of $V(G)$, called *image restrictions*, and $\mathcal{J} = \{J_x\}_{x \in V(H)}$ be a collection of subsets of $V(\Gamma) \setminus V(G)$, called *restricting vertices*. We say that $\mathcal{I}$ and $\mathcal{J}$ are a $(\rho, \zeta, \Delta, \Delta_J)$-*restriction pair* if the following properties hold for each $i \in [r]$ and $x \in X_i$.

(*a*) The set $X_i^* \subseteq X_i$ of *image restricted* vertices in X_i, that is, vertices such that $I_x \neq V_i$, has size $|X_i^*| \leq \rho |X_i|$.

(*b*) If $x \in X_i^*$, then $I_x \subseteq N_\Gamma^*(J_x; V_i)$ is of size at least $\zeta(dp)^{|J_x|} |V_i|$.

(*c*) If $x \in X_i^*$, then $|J_x| + \deg_H(x) \leq \Delta$ and if $x \notin X_i^*$, then $J_x = \emptyset$.

(*d*) Each Γ-vertex appears in at most Δ_J of the sets of $\mathcal{J}$.

(*e*) We have $\left| N_\Gamma^*(J_x; V_i) \right| = (p \pm \varepsilon p)^{|J_x|} |V_i|$.

(*f*) If $x \in X_i^*$, for each $xy \in E(H)$ with $y \in X_j$, the pair $\left(N_\Gamma^*(J_x; V_i), N_\Gamma^*(J_y; V_j) \right)$ is (ε, d, p)-regular in G.

This definition does indeed generalise the dense image restrictions of [56], since (*a*) is one of the conditions of [56], while if $p = 1$ and so Γ is the complete graph, (*b*) reduces to the other condition of [56]. We can set $J_x = \emptyset$ for all x, so that (*c*)–(*e*) become trivial, and (*f*) follows since $\mathcal{X}$ is an R-partition and $\mathcal{V}$ is an (ε, d, p)-regular R-partition. In the sparse setting the conditions amount to requiring that pre-embedded vertices creating image restrictions are embedded on 'typical vertices'.

We can now formulate our blow-up lemma for random graphs.

Lemma 1.21 (Blow-up lemma for $G(n, p)$). *For all $\Delta \geq 2$, $\Delta_{R'}$, Δ_J, $\alpha, \zeta, d > 0$, $\kappa > 1$ there exist $\varepsilon, \rho > 0$ such that for all r_1 there is a C such that for*

$$p > C \left(\frac{\log n}{n} \right)^{1/\Delta}$$

the random graph $\Gamma = G(n, p)$ a.a.s. satisfies the following.

Let R be a graph on $r \leq r_1$ vertices and let $R' \subseteq R$ be a spanning subgraph with $\Delta(R') \leq \Delta_{R'}$. Let H and $G \subseteq \Gamma$ be graphs with κ-balanced size-compatible vertex partitions $\mathcal{X} = \{X_i\}_{i \in [r]}$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$, respectively, which have parts of size at least $m \geq n/(\kappa r_1)$. Let $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ be a family of subsets of $V(H)$, let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a family of image restrictions, and $\mathcal{J} = \{J_x\}_{x \in V(H)}$ be a family of restricting vertices. Suppose that

(*BUL* 1) $\Delta(H) \leq \Delta$, $(H, \mathcal{X})$ *is an R-partition, and $\tilde{\mathcal{X}}$ is an (α, R')-buffer for $(H, \mathcal{X})$,*

(*BUL* 2) $(G, \mathcal{V})$ *is an (ε, d, p)-regular R-partition, which is (ε, d, p)-super-regular on R', has one-sided inheritance on R', and two-sided inheritance on R' for $\tilde{\mathcal{X}}$,*

(*BUL* 3) $\mathcal{I}$ *and $\mathcal{J}$ form a $(\rho, \zeta, \Delta, \Delta_J)$-restriction pair.*

Then there is an embedding $\psi \colon V(H) \to V(G)$ such that $\psi(x) \in I_x$ for each $x \in H$.

1.3.3   Degenerate graphs

We next present a version of our blow-up lemma for random graphs which allows for smaller edge probabilities in $G(n,p)$ if we want to embed graphs H whose maximum degree is much larger than their degeneracy. The *degeneracy* $\mathrm{degen}(H)$ of a graph H is the smallest integer ℓ such that H is ℓ-*degenerate*, that is, each induced subgraph of H has minimum degree at most ℓ. Equivalently, there is an order of $V(H)$ such that each vertex has at most ℓ neighbours before that vertex in the order. For example, trees have degeneracy 1 and planar graphs have degeneracy at most 5. We remark that (a variant of) the degeneracy determines the exponent in the probability p, but we nevertheless require the maximum degree of H to be bounded by a constant Δ in the lemma below. The constant in the probability p depends on Δ. For this blow-up lemma we use the same notion of regularity, super-regularity, reduced graphs, inheritance, buffer sets and image restrictions as for Lemma 1.21.

Given an order τ on $V(H)$ and a family $\mathcal{J}$ of image restricting vertices, we define

$$\pi^\tau(x) := |J_x| + \big|\{y \in N_H(x) \colon \tau(y) < \tau(x)\}\big|.$$

In other words, $\pi^\tau(x)$ denotes the number of neighbours of x that come before x, including the restricting vertices J_x.

Quantifying the dependence on τ of the probability we can work with is somewhat complex. Firstly, for a vertex x we distinguish whether it has neighbours y succeeding x in the order τ, in which case we will need to maintain one-sided inheritance properties when embedding x, or even two such neighbours y and z which form an edge, in which case we may need to maintain two-sided inheritance. Secondly, we put stricter requirements on x if it has some potential buffer vertices in its neighbourhood (since we need to maintain certain properties to embed the buffer vertices; see (*ORD* 1) and (*ORD* 3)). Thirdly, we put even stricter requirements on vertices x which are image restricted or have preceding neighbours 'far away' from x (see (*ORD* 2)). Since this is a very restrictive condition however, we allow a very small set X^e of exceptional vertices which are exempted from this rule.

Definition 1.22 ((D,p,m)-bounded order)**.** Let H be a graph given with buffer sets $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i\in[r]}$ and a restriction pair $\mathcal{I} = \{I_x\}_{x\in V(H)}$ and $\mathcal{J} = \{J_x\}_{x\in V(H)}$. Let $\tilde{X} = \bigcup_{i\in[r]}\tilde{X}_i$. Let τ be an ordering of $V(H)$ and $X^e \subseteq V(H)$. Then τ is a (D,p,m)-*bounded order* for H, $\tilde{\mathcal{X}}$, $\mathcal{I}$ and $\mathcal{J}$ with *exceptional set* X^e if the following conditions are satisfied for each $x \in V(H)$.

(*ORD* 1) Define

$$D_x := \begin{cases} D-2 & \text{if there is } yz \in E(H) \text{ with } y,z \in N_H(x) \text{ and } \tau(y),\tau(z) > \tau(x) \\ D-1 & \text{else if there is } y \in N_H(x) \text{ with } \tau(y) > \tau(x) \\ D & \text{otherwise}. \end{cases}$$

We have $\pi^\tau(x) \le D_x$, and if $x \in N(\tilde{X})$ even $\pi^\tau(x) \le D_x - 1$. Finally, if $x \in \tilde{X}$ we have $\deg(x) \le D$.

(*ORD* 2) One of the following holds:
- $x \in X^e$,
- $\pi^\tau(x) \le \frac{1}{2}D$,
- x is not image restricted and every neighbour y of x with $\tau(y) < \tau(x)$ satisfies $\tau(x) - \tau(y) \le p^{\pi^\tau(x)}m$.

(*ORD* 3) If $x \in N(\tilde{X})$ then all but at most $D - 1 - \max_{z \notin X^e} \pi^\tau(z)$ neighbours y of x with $\tau(y) < \tau(x)$ satisfy $\tau(x) - \tau(y) \leq p^D m$.

In order to obtain the best possible value of p, our aim is always to find an order τ which is (D, p, m)-bounded and minimises the value of D. To give some intuition about what typically is possible, we refer to some of the results of Section 1.2.

We can typically obtain $D \leq 2 \operatorname{degen}(H) + 1$, taking τ to be a degeneracy order for H. This gives us $\pi^\tau(x) \leq \operatorname{degen}(x)$, so that (*ORD* 2) and (*ORD* 3) hold trivially (the former with room to spare, the latter not). Furthermore, most of (*ORD* 1) is trivially satisfied (though one needs to observe that if $\operatorname{degen}(H) = 1$ then H contains no triangle). The only point which is unclear is the restriction $\deg(x) \leq D$ for $x \in \tilde{X}$. In practice (see for example the proofs of Theorems 1.1 and 1.6) one often can obtain this. The reason is that any graph H contains many vertices of degree at most $2 \operatorname{degen}(H)$, and typically we can find a partition of H in which these vertices are well-distributed among the parts.

In the event that we do not require a spanning embedding, but can afford to leave a small fraction of vertices in each part uncovered, we can obtain the slightly stronger $D \leq 2 \operatorname{degen}(H)$, as in Theorems 1.3 and 1.6. The reason for this is that we can 'pad' an almost-spanning H by adding isolated vertices to obtain a spanning H' to which we apply Lemma 1.23 (see the proof of Theorem 1.3 for details). We again use a degeneracy order τ, but this time let $\tilde{X}$ be the isolated vertices.

When we have a degeneracy order τ with the extra (bandwidth-type) property that all edges go between vertices very close together in the order, we can typically choose $D \leq \operatorname{degen}(H) + 3$. The reason is that in this case (*ORD* 2) and (*ORD* 3) are automatically satisfied, and we only need to worry about (*ORD* 1). Again, we need to be able to choose potential buffer vertices of degree at most $\operatorname{degen}(H) + 3$, but in applications this is usually possible. This applies, for example, in the case that H is an F-factor.

So far we did not mention the set X^e, or image restrictions. In applications often only very few vertices need to be image restricted, and we can typically put all of them in X^e. This is, for example, critical to obtaining a resilience result for bounded-degree trees with $p = C\left(\frac{\log n}{n}\right)^{1/3}$ in [3].

In general, the exceptional set X^e gives us the possibility to specify a small set of vertices in H for which the value of π^τ does not have to be bounded by $D/2$. This is for example useful when a few vertices have been embedded 'by hand' before the use of the blow-up lemma, creating image restrictions. Since we cannot usually select these vertices according to the degeneracy order of H, a few image restricted vertices will have more neighbours embedded 'by hand' than if H was embedded in the degeneracy order. These vertices should then be put in X^e.

Moreover, in (*ORD* 3) for each buffer neighbour x we allow a few (depending on D and τ) exceptions to the rule that embedded neighbours of x need to be close to x in the order. This is useful because vertices y of exceptionally high degree often have to come relatively early in the order to satisfy (*ORD* 1) and (*ORD* 2), and hence cannot necessarily be close to x.

We can now state our blow-up lemma for degenerate graphs. The only difference to the previous blow-up lemma is that we ask for a $(D, p, \varepsilon n/r_1)$-bounded order of H, that we allow only vertices with degree at most D as potential buffer vertices, and that the exponent in the bound on p is determined by D. We remark that for some graphs H, for example ℓ-regular graphs, which have degeneracy ℓ, this bound on p is worse than the bound in Lemma 1.21. However, for trees or planar graphs it is often much better.

Lemma 1.23 (Blow-up lemma to embed degenerate graphs in $G(n,p)$)**.**
For any $\Delta \geq 2$, $\Delta_{R'}$, Δ_J, D, $\alpha, \zeta, d > 0$, $\kappa > 1$ there exist $\varepsilon, \rho > 0$ such that for all r_1 there is a C such that for

$$p \geq C \left(\frac{\log n}{n} \right)^{1/D}$$

the random graph $\Gamma = G(n,p)$ a.a.s. satisfies the following.

    Let R be a graph on $r \leq r_1$ vertices and let $R' \subseteq R$ be a spanning subgraph with $\Delta(R') \leq \Delta_{R'}$. Let H and $G \subseteq \Gamma$ be graphs with κ-balanced, size-compatible vertex partitions $\mathcal{X} = \{X_i\}_{i \in [r]}$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$, respectively, which have parts of size at least $m \geq n/(\kappa r_1)$. Let $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ be a family of subsets of $V(H)$, let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a family of image restrictions, and $\mathcal{J} = \{J_x\}_{x \in V(H)}$ be a family of restricting vertices. Let τ be an order of $V(H)$ and $X^e \subseteq V(H)$ be a set of size $|X^e| \leq \varepsilon p^{\max_{x \in X^e} \pi^\tau(x)} n/r_1$. Suppose that

(DBUL 1) $\Delta(H) \leq \Delta$, $(H, \mathcal{X})$ is an R-partition, $\tilde{\mathcal{X}}$ is an (α, R')-buffer for $(H, \mathcal{X})$,

(DBUL 2) $(G, \mathcal{V})$ is an (ε, d, p)-regular R-partition, which is (ε, d, p)-super-regular on R', has one-sided inheritance on R', and two-sided inheritance on R' for $\tilde{\mathcal{X}}$,

(DBUL 3) $\mathcal{I}$ and $\mathcal{J}$ form a $(\rho, \zeta, \Delta, \Delta_J)$-restriction pair.

(DBUL 4) τ is a $(D, p, \varepsilon n/r_1)$-bounded order for H, $\tilde{\mathcal{X}}$, $\mathcal{I}$, $\mathcal{J}$ with exceptional set X^e.

Then there is an embedding $\psi \colon V(H) \to V(G)$ such that $\psi(x) \in I_x$ for each $x \in H$.

    Let us briefly indicate how this lemma performs in practice. As we see in for example Theorem 1.6, we can use it to obtain spanning embeddings in some situations with $D = 2\deg\mathrm{en}(H) + 1$, thus with $p \geq C \left(\frac{\log n}{n} \right)^{1/(2\deg\mathrm{en}(H)+1)}$. In [3] we obtain the promised version of Theorem 1.10 for degenerate graphs, using the same value of D. However, to obtain this result, we have to pre-embed some vertices of H, before applying Lemma 1.23, and we thus have to make some alterations to the degeneracy order to obtain an order τ. The neighbours of the pre-embedded vertices then threaten to destroy (D, p, m)-boundedness, and it is critical that we can put them into X^e.

    As a second example, if the degeneracy order (or something close to it) on H happens to have the property that all edges go only a short distance, then conditions (ORD 3) and (ORD 2) become trivially true, and we can obtain much stronger results. For example, in [3] we give a resilience result for F-factors in $G(n,p)$ with $p \geq C \left(\frac{\log n}{n} \right)^{1/(\deg\mathrm{en}(F)+3)}$, that is, using Lemma 1.23 with $D = \deg\mathrm{en}(F) + 3$. In this application we do not need any exceptions, and can set $X^e = \emptyset$.

1.3.4 Bijumbled graphs

Finally, we provide a blow-up lemma for embedding bounded degree graphs into subgraphs of sufficiently bijumbled graphs. As indicated earlier, in bijumbled graphs we use a stronger notion of regularity, which in addition to the lower bound also requires an upper bound on the edge density of subpairs.

Definition 1.24 (Regularity in bijumbled graphs)**.** In bijumbled graphs, we say that (X, Y) is (ε, d, p)-*regular* if there is a $d' \geq d$ such that for any $X' \subseteq X$ with $|X'| \geq \varepsilon|X|$ and $Y' \subseteq Y$ with $|Y'| \geq \varepsilon|Y|$, we have $d_p(X', Y') = d' \pm \varepsilon$. When we want to be it clear that we are working with this regularity concept we also call such a pair (ε, d, p)-*fully-regular*.

The reason why we use lower-regularity in random graphs is that the regularity inheritance we use in random graphs (see Section 2.1.3) provides only lower-regular pairs. On the other hand, we use full-regularity in bijumbled graphs because the regularity inheritance we prove for bijumbled graphs (again, see Section 2.1.3) requires fully-regular pairs. Unfortunately, we do not know, in either case, how to prove a regularity inheritance statement which works with the 'other' version of regularity.

All other parts of our proofs work for both lower-regularity and full-regularity. Since we would like to use many of these parts for random graphs as well as for bijumbled graphs, we let regular pairs mean lower-regular pairs whenever we work in random graphs, and fully-regular pairs whenever we work in bijumbled graphs.

In particular, super-regularity, reduced graphs, inheritance, buffer sets and image restrictions for bijumbled graphs are defined exactly as for random graphs, once the regularity concept there is replaced with the regularity concept defined here. Our blow-up lemma for bijumbled graphs then has analogous requirements and conclusions as Lemma 1.21, with the only exception that the image restrictions allowed here are much weaker: if H has maximum degree Δ we can only image restrict about a p^Δ-fraction of the vertices in any given partition class of H, rather than a small constant fraction.

Lemma 1.25 (Blow-up Lemma for bijumbled graphs). *For all $\Delta \geq 2$, $\Delta_{R'}$, Δ_J, $\alpha, \zeta, d > 0$, $\kappa > 1$ there exist $\varepsilon, \rho > 0$ such that for all r_1 there is a $c > 0$ such that if $p > 0$ and*

$$\beta \leq c p^{\max(4,(3\Delta+1)/2)} n$$

any (p, β)-bijumbled graph Γ on n vertices satisfies the following.

Let R be a graph on $r \leq r_1$ vertices and let $R' \subseteq R$ be a spanning subgraph with $\Delta(R') \leq \Delta_{R'}$. Let H and $G \subseteq \Gamma$ be graphs given with κ-balanced, size-compatible vertex partitions $\mathcal{X} = \{X_i\}_{i \in [r]}$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$, respectively, which have parts of size at least $m \geq n/(\kappa r_1)$. Let $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ be a family of subsets of $V(H)$, let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a family of image restrictions, and $\mathcal{J} = \{J_x\}_{x \in V(H)}$ be a family of restricting vertices. Suppose that

(JBUL 1) $\Delta(H) \leq \Delta$, $(H, \mathcal{X})$ is an R-partition, $\tilde{\mathcal{X}}$ is an (α, R')-buffer for $(H, \mathcal{X})$,

(JBUL 2) $(G, \mathcal{V})$ is an (ε, d, p)-regular R-partition, which is (ε, d, p)-super-regular on R', and has one-sided inheritance on R', and two-sided inheritance on R' for $\tilde{\mathcal{X}}$,

(JBUL 3) $\mathcal{I}$ and $\mathcal{J}$ form a $(\rho p^\Delta, \zeta, \Delta, \Delta_J)$-restriction pair.

Then there is an embedding $\psi \colon V(H) \to V(G)$ such that $\psi(x) \in I_x$ for each $x \in H$.

We remark that we do not believe that the bound on β in this result is even close to being best possible (this is discussed further in the concluding remarks, Section 7.2.2). However, this blow-up lemma is the first general embedding result which allows the embedding of spanning structures in bijumbled graphs (see for example Theorem 1.2). Previously such embedding results were known only for some special subgraphs H, such as Hamilton cycles [64], triangle factors [65], or powers of Hamilton cycles [5]. Moreover, our blow-up lemma allows for resilience results in pseudorandom graphs. Previously, the only results for large subgraphs in this direction that we are aware of deal with cycles. Specifically, Sudakov and Vu [75] found the local resilience of (n, d, λ)-graphs with respect to Hamiltonicity, Dellamonica, Kohayakawa, Marciniszyn, and Steger [34] found the local resilience of bijumbled graphs with respect to containing long cycles, and Krivelevich, Lee and Sudakov [62], with a

stronger bijumbledness requirement, found the local resilience with respect to pancyclicity (containing cycles of all lengths).

1.3.5 Inheritance of regularity

The final results we would like to highlight are the regularity inheritance lemmas for random graphs we mentioned earlier. These are necessary not only in many applications of our sparse blow-up lemmas, such as [3], but are also useful in other random graph contexts, for example in [4, 7]

Both statements rely crucially on the regularity inheritance work of Gerke, Kohayakawa, Rödl and Steger [42], and also use ideas from [54]. The first lemma gives an upper bound on the number of vertices which fail to have one-sided lower-regularity inheritance.

Lemma 1.26 (One-sided lower-regularity inheritance in $G(n,p)$)**.** *For each $\varepsilon', d > 0$ there are $\varepsilon_0 > 0$ and C such that for all $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, a.a.s. $\Gamma = G(n,p)$ has the following property. Let $G \subseteq \Gamma$ be a graph and X, Y be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-lower-regular in G and*

$$|X| \geq C \max\left(p^{-2}, p^{-1}\log n\right) \quad and \quad |Y| \geq Cp^{-1}\log(en/|X|),$$

then for at most $Cp^{-1}\log(en/|X|)$ vertices $z \in V(\Gamma)$ the pair $\left(N_\Gamma(z;X), Y\right)$ is not (ε', d, p)-lower-regular in G.

The corresponding two-sided inheritance lemma is the following.

Lemma 1.27 (Two-sided lower-regularity inheritance in $G(n,p)$)**.** *For each $\varepsilon', d > 0$ there are $\varepsilon_0 > 0$ and C such that for all $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, a.a.s. $\Gamma = G(n,p)$ has the following property. Let $G \subseteq \Gamma$ be a graph and X, Y be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-lower-regular in G and*

$$|Y| \geq |X| \geq C \max\left(p^{-2}, p^{-1}\log n\right),$$

then there are at most $C \max\left(p^{-2}, p^{-1}\log(en/|X|)\right)$ vertices $z \in V(\Gamma)$ such that $\left(N_\Gamma(z;X), N_\Gamma(z;Y)\right)$ is not (ε', d, p)-lower-regular in G.

These two results are quite similar to [54, Proposition 15]. The crucial difference are the bounds on the sizes of X and Y as well as the bound on the number of vertices z that do not preserve regularity. We remark that the bounds given in [54, Proposition 15] are (roughly) equivalent to our bounds when p is as small as possible in our blow-up lemmas, that is $p = \Theta\left((\log n/n)^{1/\Delta}\right)$, but for bigger p our results are stronger.

We would like to stress that these results give sharp bounds, up to the constant C, on the number of vertices which may fail to inherit lower-regularity. We prove both theorems, and that they are sharp, in Section 2.1.3.

Finally, we note that corresponding inheritance lemmas (which take as input and give as output fully-regular pairs) for bijumbled graphs were first proved by Conlon, Fox and Zhao [31]. The improved versions of these lemmas which we state in Section 2.1.3 were proved in [8].

1.4 Proof overview

The following is a high-level overview of the proofs of our blow-up lemmas. In these proofs we merge ideas that were also used in proofs of dense blow-up lemmas [26, 56, 71] and of the sparse embedding lemma in [54] with many new ingredients. Of course there are differences between the proofs of Lemmas 1.21, 1.23, and 1.25 but at the level of this overview we will avoid mentioning most of them.

In each of our blow-up lemmas we want to embed a graph H, which is given with a partition $\mathcal{X}$ and potential buffer sets $\tilde{\mathcal{X}}$, into a graph $G \subseteq \Gamma$ with a compatible super-regular partition $\mathcal{V}$, possibly with some image restrictions, and where Γ is a sparse random or bijumbled graph. To avoid technical details, we will in this overview largely ignore the image restrictions; these turn out not to play a large rôle in the proofs. Our embedding strategy for H is comprised of different randomised embedding procedures. Before we can apply them though, we need to prepare the graphs H and G by subpartitioning their partition classes suitably (details of what we need can be found in Sections 2.3.2 and 2.3.3).

We start by subpartitioning the partition classes of H such that any pair of vertices in a part of the new partition is at distance at least ten, which is possible by a trick first used by Alon and Füredi [16], and in a blow-up lemma setting by Rödl and Ruciński [71], that relies on the Hajnal-Szemerédi theorem (Theorem 2.3). Having only distant vertices in a part will provide sufficient independence of these vertices in our randomised embedding procedures. To get a compatible new partition of G, we subpartition the clusters of G randomly.

Next, we subpartition each cluster V_i of G randomly into several parts

$$V_i = V_i^{\mathrm{main}} \,\dot\cup\, V_i^{\mathrm{q}} \,\dot\cup\, V_i^{\mathrm{c}} \,\dot\cup\, V_i^{\mathrm{buf}}.$$

The first of these sets is large (of size $(1-3\mu)|V_i|$ for some small μ), while the remaining three are much smaller (of size $\mu|V_i|$). Because the subpartitionings were performed randomly, subparts of super-regular pairs (given by R') maintain super-regularity. The reason for doing this is that we perform the embedding in stages, and these stages require separate parts. This idea is also used in [26].

We also partition X_i into a large part X_i^{main} (of size at most $(1-4\mu)|X_i|$) and a small part X_i^{buf} (of size $4\mu|X_i|$), called the set of *buffer vertices*. The latter set is required to be a subset of the potential buffer vertices $\tilde{X}_i$ in X_i. We also require various extra properties of X_i^{buf}. In the proof of the random graphs blow-up lemma, Lemma 1.21, we also include another small subpart X_i^{c} in this partitioning step. This is due to the fact that in this proof we have to treat the case that most of the vertices $\tilde{X}_i$ are contained in copies of $K_{\Delta+1}$ specially (in order to obtain the claimed bound on the probability p). We need the sets V_i^{c} to embed this small subpart.

At this point H and G are prepared for the embedding. We now describe our randomised embedding approach. Firstly, we make use of a *random greedy algorithm* (RGA) to embed X_i^{main} into V_i^{main} (see Section 3.2 for the simplest version used in the proof of Lemma 1.21). An algorithm of this type was also used by Komlós, Sárközy and Szemerédi [56] to prove the dense blow-up lemma. The RGA embeds the sets X_i^{main} vertex by vertex, in each step avoiding some *bad* vertices of G and embedding the current vertex $x \in X_i^{\mathrm{main}}$ into its so-called *candidate set*. If ψ is the *partial embedding* of H into G that we have constructed so far, then the candidate set $C(x) \subseteq V_i$ is the set of vertices adjacent in G to each $\psi(y)$ such that y is an embedded neighbour of x in H. However, some of the vertices in $C(x)$ may have been used as images for other vertices already, so we let $A(x) = C(x) \setminus \mathrm{Im}(\psi)$ be the *available candidate set* for x. Obviously, x has to be embedded into $A(x)$ in order to obtain an embedding.

In order to succeed with this strategy, we have to maintain certain properties for the partial embedding ψ, which we call *good partial embedding* properties (see Section 2.3.5). We remark that the properties of a restriction pair ensure that the trivial partial embedding, which we have at the beginning when no vertices of H are embedded yet, is a good partial embedding. This fact means that in the proofs we usually do not need to distinguish between vertices which are and are not image restricted (and it also justifies that we ignore image restrictions in this overview).

Among other properties, such as regularity properties, in a good partial embedding we require that for each unembedded $x \in X_i$ the set $C(x)$ is large, that is of size $\Omega(d^\ell p^\ell |V_i^{\text{main}}|)$ where ℓ is the number of already embedded neighbours of x. In order to maintain this property we need to avoid certain bad vertices $B(x)$ when embedding x. More precisely, when we embed x this leads to a change of the candidate sets $C(y)$ of unembedded neighbours y of x. The set $B(x)$ contains the vertices $v \in V_i$ that would lead to some $C(y)$ becoming small (or have some other bad properties that would prevent us from maintaining a good partial embedding; see Section 2.3.6). The RGA then embeds x *uniformly at random* into $V_i^{\text{main}} \cap A(x) \setminus B(x)$.

That we choose the images randomly helps us in several ways. Most immediately, we can use it to show that the sets $C(y)$ for unembedded vertices $y \in X_i^{\text{main}} \cup X_i^{\text{buf}}$ are 'uniformly' distributed over V_i^{main}. In particular, this means that they tend not to be contained entirely in $\text{Im}(\psi)$, which, together with the fact that $B(y)$ is always small, implies that the set $V_i^{\text{main}} \cap A(y) \setminus B(y)$ to which we wish to embed a future $y \in X_i^{\text{main}}$ is usually not small.

Unfortunately though, the RGA will not succeed in embedding every vertex of X_i^{main}, because occasionally we will come across a vertex x such that $V_i^{\text{main}} \cap A(x) \setminus B(x)$ is small. We put such vertices x in a *queue* $X_i^{\text{q}} \subseteq X_i^{\text{main}}$. The use of such a queue appears already in [56]. The 'uniform' distribution of the sets $C(x)$ allows us to show that X_i^{q} remains much smaller than the so far untouched V_i^{q}. We will show that this allows us to embed X_i^{q} into V_i^{q}, maintaining a good partial embedding. We note that in the proof of the random graphs blow-up lemma, Lemma 1.21, we perform this embedding after we have embedded all other vertices of X_i^{main}, using a matching strategy similar to that of Kohayakawa, Rödl, Schacht and Szemerédi [54]. In the proofs of the other two blow-up lemmas, Lemma 1.23 and 1.25, however, we embed the vertices of X_i^{main} in the given order, again using a random greedy strategy. In either case, the underlying idea is to show that the only reason why we might fail to embed X_i^{q} into V_i^{q} is that G contains a 'dense spot', contradicting the fact that our random or bijumbled Γ does not contain such dense spots.

At this point, all vertices of X_i^{main} are embedded, and it remains to embed the carefully chosen buffer vertices X_i^{buf}. Since any two buffer vertices are at distance at least 10 in H, in particular all neighbours of buffer vertices have been embedded, so the candidate set $C(x)$ of a buffer vertex x will not change anymore, and it suffices to find a system of distinct representatives for the available candidate sets $A(x)$, which we achieve by verifying Hall's condition. Thus, for each set $Y \subseteq X_i^{\text{buf}}$, we need to show that the union W of the available candidate sets for $y \in Y$ satisfies $|W| \geq |Y|$. We separate three cases: Y is either a small fraction of X_i^{buf}, or most of X_i^{buf}, or somewhere intermediate. In the first case we use essentially the same argument as when embedding the queue vertices. In the last case we show that the 'uniform' distribution of the sets $C(x)$ for $x \in X_i^{\text{buf}}$ implies that W is most of $V_i \setminus \text{Im}(\psi)$.

To handle the remaining case when Y contains most or all of X_i^{buf}, we show that the RGA gives us an extra property: for every vertex $v \in V_i$ there are many vertices $x \in X_i^{\text{buf}}$ such that v is a candidate for x. This is the point in the proof where we use the super-regularity and inheritance properties that G

satisfies on the subgraph R' of R. We also remark that in the case of the random graphs blow-up lemma, Lemma 1.21, we cannot in general establish this property: When most of the vertices in $\tilde{X}_i$ are contained in copies of $K_{\Delta+1}$ we obtain the claimed feature only for most vertices instead of all vertices $v \in V_i$. Hence in this case, to recover the special property, we perform an additional embedding stage, using the sets X_i^c and V_i^c, to fix these *buffer defects* (see Section 3.4). This additional embedding stage also uses the super-regularity and inheritance properties of R'. We do not use these properties elsewhere in the proof. Finally, using again that Γ and hence G does not have 'dense spots' we can show that the above described extra property implies $|W| \geq |Y|$ as desired.

Summarising, our blow-up lemma proofs contain three main embedding procedures: the random greedy algorithm, the queue embedding, and the embedding of the buffer vertices. In the proof of Lemma 1.21 we also perform an additional embedding procedure for fixing buffer defects.

One important difference between the proofs of Lemmas 1.21 and 1.25, and of Lemma 1.23, is that in the former we choose the order of embedding vertices in the RGA, whereas in the latter an order is given to us. We change it only by moving the vertices X_i^{buf} to the end of the order.

Chapter 2

Tools, preparation, and setup

In this chapter we lay the foundations for the proofs of our blow-up lemmas. We start in Section 2.1 by providing various tools we need: a collection of concentration inequalities, a suitable variant of the Hajnal-Szemerédi theorem on equitable partitions of graphs, and regularity inheritance lemmas in random and in bijumbled graphs. In Section 2.2 we then establish various deterministic properties that the random graph or a bijumbled graph enjoys and which we use in our proofs. In Section 2.3 we describe the general setup that is used in the proofs of our three blow-up lemmas. In Section 2.4 finally we collect a series of lemmas which are useful for analysing the various random greedy algorithms our proofs use.

2.1 Tools

2.1.1 Probabilistic inequalities

We need the following forms of Chernoff's inequality and of the hypergeometric inequality. Recall that if X is hypergeometrically distributed with parameters N, n and m, then $\mathbb{E}X = mn/N$.

Theorem 2.1 (Corollary 2.4, Theorems 2.8 and 2.10 from [46])**.** *Suppose X is a random variable which is either the sum of a collection of independent Bernoulli random variables, or is hypergeometrically distributed. Then we have for $\delta \in (0, 3/2)$*

$$\mathbb{P}\big(X > (1+\delta)\mathbb{E}X\big) < e^{-\delta^2 \mathbb{E}X/3} \quad \text{and} \quad \mathbb{P}\big(X < (1-\delta)\mathbb{E}X\big) < e^{-\delta^2 \mathbb{E}X/3}\,.$$

We also have, for any $t \geq 6\mathbb{E}X$,

$$\mathbb{P}\big(X \geq \mathbb{E}X + t\big) \leq e^{-t}\,.$$

We will often want to bound above the sum of a sequence of Bernoulli random variables $Y_1, \ldots, Y_n$ coming from some process which are *not* independent, but which have the following *sequentially dependent* structure. Suppose $\mathscr{H}_0, \ldots, \mathscr{H}_n$ are increasing 'histories' of the process, that is, information on all the random choices made in the process up to given increasing times in the process. Suppose that for each $1 \leq i \leq n$, the value of the random variable Y_i is determined by $\mathscr{H}_i$, and that we have functions $p_i = p_i(\mathscr{H}_{i-1})$ such that $\mathbb{E}[Y_i|\mathscr{H}_{i-1}] \leq p_i$ holds *almost surely*, that is, with probability 1. If $\sum_{i=1}^n p_i$ is

almost surely bounded above by x, then the following lemma claims the same upper tail bound on $\sum_{i=1}^{n} Y_i$ holds as we would get from Theorem 2.1 if the Y_i were independent and the sum of their expectations were x, and also gives the same lower tail bound as Theorem 2.1 under similar conditions.

It is convenient to phrase this lemma in terms of a sequence of partitions $\mathcal{F}_0, \ldots, \mathcal{F}_n$, each refining the previous, of a probability space Ω. For the connection to processes and histories, observe that any finite stochastic process is associated with the finite probability space of all possible outcomes, with the probability measure coming from the process. The possible histories of the process up to any given time t naturally give a partition of this probability space, and two histories up to an earlier and later time give two partitions, the first refined by the second. Thus the following lemma indeed applies to sequentially dependent random variables as described above.

We note that this lemma could be phrased in terms of a filtration and random variables measurable with respect to elements of the filtration, which might be more familiar to readers with a background in probability. That we do not use this notation here is purely to avoid defining these concepts. Furthermore, we remark that the lemma is essentially a super/submartingale inequality, proved in the standard way.[1] However, we did not find this particular inequality in the literature, so give a proof from first principles here.

Lemma 2.2 (Sequential dependence lemma). *Let Ω be a finite probability space, and $\mathcal{F}_0, \ldots, \mathcal{F}_n$ be partitions of Ω, with $\mathcal{F}_{i-1}$ refined by $\mathcal{F}_i$ for each $i \in [n]$. For each $i \in [n]$ let Y_i be a Bernoulli random variable on Ω which is constant on each part of $\mathcal{F}_i$, and let p_i be a real-valued random variable on Ω which is constant on each part of $\mathcal{F}_{i-1}$. Let x be a real number, $\delta \in (0, 3/2)$, and $X = Y_1 + \cdots + Y_n$.*

(a) If $\sum_{i=1}^{n} p_i \leq x$ holds almost surely, and $\mathbb{E}[Y_i | \mathcal{F}_{i-1}] \leq p_i$ holds almost surely for all $i \in [n]$, then

$$\mathbb{P}(X > (1+\delta)x) < e^{-\delta^2 x/3}.$$

(b) If $\sum_{i=1}^{n} p_i \geq x$ holds almost surely, and $\mathbb{E}[Y_i | \mathcal{F}_{i-1}] \geq p_i$ holds almost surely for all $i \in [n]$, then

$$\mathbb{P}(X < (1-\delta)x) < e^{-\delta^2 x/3}.$$

Proof. We start with (a). We shall first show by induction on n that if $\sum_{i=1}^{n} p_i \leq x$ holds almost surely and $\mathbb{E}[Y_i | \mathcal{F}_{i-1}] \leq p_i$ holds almost surely for all $i \in [n]$, then we have the following bound on the moment generating function of X for every $u \geq 0$:

$$\mathbb{E}e^{uX} = \mathbb{E}\prod_{i=1}^{n} e^{uY_i} \leq \left(1 - \frac{x}{n} + \frac{x}{n}e^u\right)^n. \tag{2.1}$$

For $n = 1$ we have $\mathbb{P}(Y_1 = 1 | \mathcal{F}_0) \leq p_1 \leq x$ almost surely, hence $\mathbb{P}(Y_1 = 1) \leq x$ almost surely. We conclude that indeed

$$\mathbb{E}e^{uY_1} = 1 - \mathbb{P}(Y_1 = 1) + e^u \mathbb{P}(Y_1 = 1) \leq 1 - x + xe^u,$$

because $1 - t + te^u$ is non-decreasing in t.

For $n \geq 2$, we shall use for each $F \in \mathcal{F}_1$ with $\mathbb{P}(F) > 0$, the induction hypothesis applied to the $n-1$ random variables $p_2|F, \ldots, p_n|F$ and the $n-1$ random variables $Y_2|F, \ldots, Y_n|F$. This is possible because,

---

[1]We would like to thank Oliver Riordan and Ori Gurel-Gurevich for pointing this out to us.

since $\mathbb{P}(F) > 0$, we have $\sum_{i=2}^{n} p_i|F \leq x - p_1|F$ almost surely and $\mathbb{E}[Y_i|F, \mathcal{F}_{i-1}] \leq p_i|F$ almost surely for each $2 \leq i \leq n$. We conclude by induction that

$$
\begin{aligned}
\mathbb{E}\prod_{i=1}^{n} e^{uY_i} &= \sum_{F \in \mathcal{F}_1} \mathbb{P}(F)e^{uY_1|F}\mathbb{E}\Big[\prod_{i=2}^{n} e^{uY_i}\Big|F\Big] \\
&\leq \sum_{F \in \mathcal{F}_1} \mathbb{P}(F)e^{uY_1|F}\left(1 - \frac{x - p_1|F}{n-1} + \frac{x - p_1|F}{n-1}e^u\right)^{n-1}.
\end{aligned}
\tag{2.2}
$$

Further, because Y_1 is constant on each part F of $\mathcal{F}_1$ we have for each $F_0 \in \mathcal{F}_0$ that

$$
\begin{aligned}
\sum_{F \in \mathcal{F}_1, F \subseteq F_0} \mathbb{P}(F)e^{uY_1|F} &= \mathbb{P}(F_0)\mathbb{E}[e^{uY_1}|F_0] \\
&= \mathbb{P}(F_0)\big(1 - \mathbb{P}(Y_1 = 1|F_0) + e^u\mathbb{P}(Y_1 = 1|F_0)\big) \leq \mathbb{P}(F_0)(1 - p_1|F_0 + e^u p_1|F_0),
\end{aligned}
$$

where the inequality uses that $1 - t + te^u$ is non-decreasing in t. Because p_1 is constant on each part F_0 of $\mathcal{F}_0$ it thus follows from (2.2) that

$$
\begin{aligned}
\mathbb{E}\prod_{i=1}^{n} e^{uY_i} &\leq \sum_{F_0 \in \mathcal{F}_0} \sum_{F \in \mathcal{F}_1, F \subseteq F_0} \mathbb{P}(F)e^{uY_1|F}\left(1 - \frac{x - p_1|F_0}{n-1} + \frac{x - p_1|F_0}{n-1}e^u\right)^{n-1} \\
&\leq \sum_{F_0 \in \mathcal{F}_0} \mathbb{P}(F_0)\big(1 - p_1|F_0 + e^u p_1|F_0\big)\left(1 - \frac{x - p_1|F_0}{n-1} + \frac{x - p_1|F_0}{n-1}e^u\right)^{n-1}.
\end{aligned}
\tag{2.3}
$$

Since $f(t) = \ln(1 - t + te^u)$ is a concave function in t it follows from Jensen's inequality that

$$
f(t) + (n-1)f\left(\frac{x-t}{n-1}\right) \leq n \cdot f\left(\frac{1}{n}\left(t + (n-1)\frac{x-t}{n-1}\right)\right) = n \cdot f\left(\frac{x}{n}\right),
$$

for any $t \in [0, 1]$. Substituting this, with $t = p_1|F_0$ for each $F_0 \in \mathcal{F}_0$, into (2.3) yields (2.1), as desired.

We are now in a position to apply Bernstein's method of applying Markov's inequality to the moment generating function. Since e^{uz} is strictly increasing in z for each $u > 0$, we have

$$
\begin{aligned}
\mathbb{P}\big(X > (1+\delta)x\big) = \mathbb{P}\big(e^{uX} > e^{(1+\delta)ux}\big) &\leq \big(\mathbb{E}e^{uX}\big)e^{-(1+\delta)ux} \\
&\overset{(2.1)}{\leq} \big(1 - \tfrac{x}{n} + \tfrac{x}{n}e^u\big)^n e^{-(1+\delta)ux},
\end{aligned}
\tag{2.4}
$$

if $u > 0$, where the first inequality is Markov's inequality. We may assume $(1+\delta)x < n$, since otherwise $\mathbb{P}\big(X > (1+\delta)x\big) = 0$ and the lemma statement is trivial. Thus we can choose u such that $e^u = \frac{(1+\delta)(n-x)}{n-(1+\delta)x}$, plug that into the right hand side of (2.4) and obtain

$$
\begin{aligned}
\mathbb{P}\big(X > (1+\delta)x\big) &\leq \left(1 + \frac{\delta x}{n-(1+\delta)x}\right)^n \left(\frac{(1+\delta)(n-x)}{n-(1+\delta)x}\right)^{-(1+\delta)x} \\
&= \left(1 + \frac{\delta x}{n-(1+\delta)x}\right)^{n-(1+\delta)x} \left(\left(1 + \frac{\delta x}{n-(1+\delta)x}\right)^{-1}\left(\frac{(1+\delta)(n-x)}{n-(1+\delta)x}\right)\right)^{-(1+\delta)x} \\
&\leq e^{\delta x}(1+\delta)^{-(1+\delta)x} \leq e^{-\delta^2 x/3},
\end{aligned}
$$

where the last inequality uses $(1+\delta)\ln(1+\delta) - \delta \geq \delta^2/3$. This proves (a).

The proof of (b) is similar. By the analogous calculation, we have

$$\mathbb{E}e^{u(n-X)} \leq \left(1 - \tfrac{n-x}{n} + \tfrac{n-x}{n}e^u\right)^n,$$

and again Bernstein's method, this time choosing $e^u = \frac{n-x+\delta x}{(n-x)(1-\delta)}$, gives the desired result for $\delta \in (0,1)$:

$$
\begin{aligned}
\mathbb{P}(X < (1-\delta)x) = \mathbb{P}(n-X > n-x+\delta x) &= \mathbb{P}\left(e^{u(n-X)} > e^{u(n-x+\delta x)}\right) \\
&\leq \mathbb{E}[e^{u(n-X)}]e^{-u(n-x+\delta x)} \leq \left(1 - \tfrac{n-x}{n} + \tfrac{n-x}{n}e^u\right)^n e^{-u(n-x+\delta x)} \\
&= \left(1 + \tfrac{n-x}{n}\left(\tfrac{\delta n}{(n-x)(1-\delta)}\right)\right)^n \left(\tfrac{n-x+\delta x}{(n-x)(1-\delta)}\right)^{-(n-x+\delta x)} \\
&= (1-\delta)^{-n}\left(\tfrac{n-x+\delta x}{(n-x)(1-\delta)}\right)^{-(n-x+\delta x)} = (1-\delta)^{-(1-\delta)x}\left(\tfrac{n-x+\delta x}{n-x}\right)^{-(n-x+\delta x)} \\
&\leq (1-\delta)^{-(1-\delta)x}\left(1 + \tfrac{\delta x}{n-x}\right)^{-(n-x)} \leq (1-\delta)^{-(1-\delta)x}e^{-\delta x} \leq e^{-\delta^2 x/3}.
\end{aligned}
$$

Note that in this case we only need to consider $\delta \in (0,1)$ as the probability of $X < 0$ is zero, so this gives (b). $\qquad\square$

2.1.2 Equitable partitions

The Hajnal-Szemerédi theorem states that any graph F has a k-colouring with equitable colour classes for every $k \geq (\Delta(F)+1)$.

Theorem 2.3 (Hajnal-Szemerédi [44])**.** *Given any graph F and $k \geq \Delta(F)+1$, there is an equitable partition $V(F) = V_1 \dot\cup \cdots \dot\cup V_k$ such that each part is independent.*

For preparing the graph H for the embedding in the proofs of our blow-up lemmas, we require an equitable partitioning result similar to the Hajnal-Szemerédi theorem (which we shall apply to an auxiliary graph F defined for each part of the given partition of H). The difference is that we want to specify a subset X of the vertices of F and obtain an equitable partition of F that also equitably partitions X (this subset will be the buffer set). Since for us it is not essential to obtain a sharp bound on the number of parts required, this result is not difficult to deduce from the Hajnal-Szemerédi theorem.

Lemma 2.4. *Given any graph F and a subset X of $V(F)$, there is an equitable partition $V(F) = V_1 \dot\cup \cdots \dot\cup V_{8\Delta(F)}$ such that each part is independent and the sets $V_i \cap X$ form an equitable partition of X.*

Proof. We may assume without loss of generality that $|X| \leq v(F)/2$, since otherwise we could replace X with $V(F) \setminus X$. We now first use the Hajnal-Szemerédi theorem to get an equitable partition of $F[X]$ into $8\Delta(F)$ independent parts. We proceed by adding the remaining vertices of F to these parts, maintaining their independence, to obtain a partition $V_1 \dot\cup \cdots \dot\cup V_{8\Delta(F)}$ of $V(F)$. Among all possible such partitions we choose the one which is as equitable as possible. Observe that in this process it is always possible to put every vertex into some part: for any vertex, there are at most $\Delta(F)$ parts to which it cannot be added, so at least $7\Delta(F)$ parts exist to which it can be added to form an independent set.

We are now done if there exist no pair of parts V_i, V_j whose sizes differ by two or more. Suppose V_1 is the smallest part. Observe that there are at most $\Delta(F)|V_1|$ vertices in $V(F) \setminus X$ which have a neighbour in V_1. Since $8\Delta(F)|V_1| \leq v(F)$ it follows that there are at least $\frac{3}{8}v(F)$ vertices in $V(F) \setminus X$ which have no neighbour in V_1.

Assume for contradiction that there are parts V_i with $|V_i| \geq |V_1| + 2$. Observe that for each such i every vertex of $V_i \setminus X$ is adjacent to some vertex of V_1, or we would be able to move a vertex from $V_i \setminus X$ to V_1 and obtain a more equitable partition. Hence the at least $\frac{3}{8}v(F)$ vertices in $V(F) \setminus X$ which have no neighbour in V_1 are all contained in parts of size at most $|V_1| + 1 \leq 2|V_1|$, and thus lie in at least $\frac{3}{2}\Delta(F)$ different parts. Now fix a set V_i with $|V_i| \geq |V_1| + 2$, and let v be any vertex of $V_i \setminus X$ (which exists because we started with an equitable partition of X). Since v has at most $\Delta(F) < \frac{3}{2}\Delta(F)$ neighbours, there is a set V_k which contains no vertex adjacent to v, and which contains a vertex w, not in X, with no neighbours in V_1. We replace V_i with $V_i \setminus \{v\}$, V_k with $V_k \cup \{v\} \setminus \{w\}$, and V_1 with $V_1 \cup \{w\}$. The result is a more equitable partition, a contradiction. $\qquad\square$

2.1.3 Regularity inheritance

In our blow-up lemmas we require certain regularity inheritance properties. In this subsection we justify that we can obtain these in subgraphs of random or bijumbled graphs. More precisely, we provide inheritance lemmas stating that most vertices in a (sparse) regular partition satisfy these inheritance properties. This is important for the correctness of our embedding procedures proving the blow-up lemmas.

Recall that, according to our definition, regularity in random graphs and regularity in bijumbled graphs are different things: lower-regularity in the former and full-regularity in the later. Recall also that the sole reason for this difference is that we can only establish regularity inheritance for the respective version, that is, that the lemmas provided in this subsection require lower-regularity when they concern random graphs and full-regularity when they concern bijumbled graphs. Accordingly, in this subsection we shall make the two different regularity concepts explicit and talk about lower-regularity and full-regularity for clarity.

For subgraphs G of bijumbled graphs Γ we can simply rely on the following inheritance lemmas from [8]. These lemmas consider three disjoint vertex sets X, Y, Z such that (X, Y) forms a fully-regular pair. The first lemma is a one-sided regularity inheritance statement and states that for most vertices $z \in Z$ the pair $(N_\Gamma(z; X), Y)$ inherits full-regularity.

Lemma 2.5 (One-sided inheritance in bijumbled graphs [8, Lemma 3]).
For each $\varepsilon', d > 0$ there are $\varepsilon_0, c_0 > 0$ such that for all $0 < \varepsilon \leq \varepsilon_0$ and $0 < p < 1$ the following holds. Let $G \subseteq \Gamma$ be graphs and X, Y, Z be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-fully-regular in G and Γ is

$$\left(p, c_0 p^2 (\log_2 \tfrac{1}{p})^{-1/2} \sqrt{\max\left(|X||Y|, |X||Z|\right)} \right)\text{-bijumbled}$$

then for at most $\varepsilon'|Z|$ vertices z of Z the pair $(N_\Gamma(z; X), Y)$ is not (ε', d, p)-fully-regular in G.

The second lemma requires stronger bijumbledness for establishing two-sided regularity inheritance. It states that for most $z \in Z$ the pair $(N_\Gamma(z; X), N_\Gamma(z; Y))$ inherits full-regularity.

Lemma 2.6 (Two-sided inheritance in bijumbled graphs [8, Lemma 4]).
For each $\varepsilon', d > 0$ there are $\varepsilon_0, c_0 > 0$ such that for all $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$ the following holds. Let $G \subseteq \Gamma$ be graphs and X, Y, Z be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-fully-regular in G and Γ is

$$\left(p, c_0 p^3 \sqrt{\max\left(|X||Y|, |X||Z|, |Y||Z| \right)} \right)\text{-bijumbled}$$

then for at most $\varepsilon'|Z|$ vertices z of Z the pair $\left(N_\Gamma(z; X), N_\Gamma(z; Y) \right)$ is not (ε', d, p)-fully-regular in G.

The remainder of this subsection is devoted to proving the inheritance lemmas for random graphs, which we stated in Section 1.3.5, and showing their optimality. For the reader's convenience, we restate these lemmas here, beginning with the one-sided inheritance lemma.

Lemma 1.26. *For each $0 < \varepsilon', d$ there are $\varepsilon_0 > 0$ and C such that for all $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, a.a.s. $\Gamma = G(n, p)$ has the following property. Let $G \subseteq \Gamma$ be a graph and X, Y be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-lower-regular in G and*

$$|X| \geq C \max\left(p^{-2}, p^{-1} \log n \right) \quad and \quad |Y| \geq C p^{-1} \log(en/|X|),$$

then for at most $C p^{-1} \log(en/|X|)$ vertices $z \in V(\Gamma)$ the pair $\left(N_\Gamma(z; X), Y \right)$ is not (ε', d, p)-lower-regular in G.

The corresponding two-sided inheritance lemma is:

Lemma 1.27. *For each $0 < \varepsilon', d$ there are $\varepsilon_0 > 0$ and C such that for all $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, a.a.s. $\Gamma = G(n, p)$ has the following property. Let $G \subseteq \Gamma$ be a graph and X, Y be disjoint subsets of $V(\Gamma)$. If (X, Y) is (ε, d, p)-lower-regular in G and*

$$|Y| \geq |X| \geq C \max\left(p^{-2}, p^{-1} \log n \right),$$

then there are at most $C \max\left(p^{-2}, p^{-1} \log(en/|X|) \right)$ vertices $z \in V(\Gamma)$ such that $\left(N_\Gamma(z; X), N_\Gamma(z; Y) \right)$ is not (ε', d, p)-lower-regular in G.

Note that the regularity inheritance lemmas in random graphs, Lemmas 1.26 and 1.27 are stated in a different form to the ones for bijumbled graphs above, not requiring any bound on p, but rather having p appear in the estimates for the sizes of X and Y as well as the number of vertices not preserving regularity. We state them in this form because this seems most suitable when we want to use them in applications of our blow-up lemma. For bijumbled graphs on the other hand we decided not to switch to this form since we would obtain more complicated conditions in this case (bounding the products of set sizes such as $|X||Y|$, instead of the sets themselves).

In the proofs of Lemmas 1.26 and 1.27 we follow the approach of [54], and rely on the following result of Gerke, Kohayakawa, Rödl and Steger [42] stating that the vast majority of subpairs (X', Y) of a lower-regular pair (X, Y) in any graph inherit lower-regularity. Unfortunately this result becomes false if we replace lower-regularity by full-regularity, and it is precisely this reason why we need to work with lower-regularity in random graphs.

Theorem 2.7 (Theorem 3.6 from [42])**.** *For any d, β, $\varepsilon' > 0$ there exist $\varepsilon_0 > 0$ and C such that for any $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, if (X,Y) is an (ε,d,p)-lower-regular pair in a graph G, then the number of sets $X' \subseteq X$ with $|X'| = w \geq C/p$ such that (X',Y) is an (ε',d,p)-lower-regular pair in G is at least $(1 - \beta^w)\binom{|X|}{w}$.*

The following is an analogous statement for subpairs (X',Y') of (X,Y) and is an immediate consequence of Theorem 2.7.

Corollary 2.8 (Corollary 3.8 from [42])**.** *For any d, β, $\varepsilon' > 0$ there exist $\varepsilon_0 > 0$ and C such that for any $0 < \varepsilon < \varepsilon_0$ and $0 < p < 1$, if (X,Y) is an (ε,d,p)-lower-regular pair in a graph G, then the number of pairs $X' \subseteq X$ and $Y' \subseteq Y$ with $|X'| = w_1 \geq C/p$ and $|Y'| = w_2 \geq C/p$ such that (X',Y') is an (ε',d,p)-lower-regular pair in G is at least $(1 - \beta^{\min(w_1,w_2)})\binom{|X|}{w_1}\binom{|Y|}{w_2}$.*

The proof idea for Lemma 1.26 is the following. First, we will show that if (X,Y) is a counterexample, that is, a lower-regular pair not satisfying the conclusion of the lemma, then there are $X' \subseteq X$ and $Y' \subseteq Y$ which are only slightly smaller than X and Y respectively, and a set Z' of vertices outside $X' \cup Y'$ whose neighbourhoods in X' all have size about $p|X'|$ but which nevertheless do not form regular pairs with Y'. We then find $Y^* \subseteq Y'$ which has the same properties as Y', but which has size $\frac{1}{2}Cp^{-1}\log(en/|X|)$. We finally argue that this last structure is unlikely to exist in $G(n,p)$, completing the proof. It may not be obvious what we gain by reducing (X',Y') to (X',Y^*); what we gain is that there are not too many edges of Γ between X' and Y^*, and hence not too many possible lower-regular subgraphs of Γ. This allows us to take a union bound over all possible choices, which would fail if we attempted it with (X',Y').

The proof of Lemma 1.27 is similar but somewhat more complicated, and we will sketch it after proving Lemma 1.26. Before embarking on these proofs, we need three preliminary lemmas.

The first lemma assumes we are given a lower-regular pair together with a collection of ℓ subpairs which do not inherit lower-regularity, and asserts that we can scale down the sizes of this pair and the subpairs in many ways. For Lemma 1.26 we only need one way; for Lemma 1.27 we will need the full power of this lemma.

Lemma 2.9. *For each $d,\varepsilon_2,\delta > 0$ there exist ε_1 and C such that for all $0 < \varepsilon < \varepsilon_1$ the following holds. Let (X,Y) be an (ε,d,p)-lower-regular pair in a graph G, and suppose $|X| \geq m \geq Cp^{-1-i}$ for $i \in \{0,1\}$. Further, let $X_j \subseteq X$ and $Y_j \subseteq Y$ with $j \in [\ell]$ form ℓ subpairs of (X,Y) for some $\ell \in [m]$, with $|X_j| \geq \frac{1}{100}p^i|X|$ for each j, such that (X_j,Y_j) is not (δ,d,p)-lower-regular for any $j \in [\ell]$.*

Then there are at least $\frac{1}{2}\binom{|X|}{m}$ choices of $X^* \subseteq X$ with $|X^*| = m$ such that (X^*,Y) is (ε_2,d,p)-lower-regular, such that $|X_j \cap X^*| = \left(1 \pm \frac{1}{2}\right)m|X_j|/|X|$ for each $j \in [\ell]$, and such that for each $j \in [\ell]$, the pair $(X_j \cap X^*, Y_j)$ is not $(\delta^2/10,d,p)$-lower-regular.

Proof. Given $d,\varepsilon_2 > 0$, let ε_1 and C' be returned by Theorem 2.7 with input $d, \frac{1}{4}, \varepsilon_2$. Given $\delta > 0$, let $C \geq C'$ be such that $3(x^2/C)\exp(-\delta_1^2 x/3600) < \frac{1}{4}$ for each $x \geq 1$.

Given $0 < \varepsilon < \varepsilon_1$, and given an (ε,d,p)-lower-regular pair (X,Y) with $|X| \geq m \geq Cp^{-1-i}$, let $X^* \subseteq X$ be chosen uniformly at random from the m-sized subsets of X. By Theorem 2.7, with probability at least $\frac{3}{4}$, the pair (X^*,Y) is (ε_2,d,p)-lower-regular.

Given any pair (X_j,Y_j), since the pair is not (δ,d,p)-regular, there exist subsets X_j' and Y_j' of X_j and Y_j respectively, of sizes $\delta|X_j|$ and $\delta|Y_j|$ respectively, such that $d_p(X_j',Y_j') \leq d - \delta$. There thus exists a subset

X_j'' of X_j' consisting of $\delta^2|X_j|/3$ vertices each of whose degree into Y_j' is smaller than $(d-\delta/2)p|Y_j'|$, as otherwise we would have

$$e(X_j' \setminus X_j'', Y_j') \geq (\delta - \delta^2/3)|X_j|(d-\delta/2)p|Y_j'| > (d-\delta)p|X_j'||Y_j'| > e(X_j', Y_j'),$$

a contradiction. Let $X_j^* = X_j'' \cap X^*$, then $d_p(X_j^*, Y_j') < d - \delta/2$, by definition of X_j''. The quantity $|X_j^*|$ is hypergeometrically distributed with mean at least $\delta^2 p^i m/300$ since $|X_j''| = \delta^2|X_j|/3$ and $|X_j| \geq \frac{1}{100}p^i|X|$. By Theorem 2.1, the probability that $|X_j^*| < \delta^2|X_j|m/(6|X|)$ is at most $\exp(-\delta^2 p^i m/3600)$. Similarly, the quantity $|X_j \cap X^*|$ is hypergeometrically distributed with mean at least $\delta^2 p^i m/300$, so by Theorem 2.1 the probability that $|X^* \cap X_j| \neq \left(1 \pm \frac{1}{2}\right)|X_j|m/|X|$ is at most $2\exp(-\delta^2 p^i m/3600)$.

By choice of m, we have $(p^i m)^2 \geq Cp^{-1}p^i m \geq Cm$, so setting $x = p^i m \geq 1$, by choice of C we have $3(x^2/C)\exp(-\delta^2 x/3600) < \frac{1}{4}$. Thus taking a union bound, we conclude that with probability at least

$$1 - \frac{1}{4} - 3m\exp(-\delta^2 p^i m/3600) > \frac{1}{2}$$

none of the above bad events occurs. In particular, for at least $\frac{1}{2}\binom{|X|}{m}$ choices of X^*, for each $j \in [\ell]$ the pair $(X_j \cap X^*, Y_j)$ is not $(\delta^2/10, d, p)$-regular, giving the conclusion of the lemma. $\square$

The next easy lemma shows that adding a few vertices to an (ε, d, p)-lower-regular pair does not destroy regularity. We note that the corresponding result for an (ε, d, p)-fully-regular pair would require knowing that G does not contain dense spots.

Lemma 2.10. *Let $0 < \varepsilon < \frac{1}{10}$. Let G be a graph and let $U', V' \subseteq V(G)$ be disjoint sets such that (U', V') is (ε, d, p)-lower-regular in G. If $U \supseteq U'$ with $|U| \leq \left(1 + \frac{1}{10}\varepsilon^3\right)|U'|$ and $V \supseteq V'$ with $|V| \leq \left(1 + \frac{1}{10}\varepsilon^3\right)|V'|$ are disjoint, then (U, V) is $(2\varepsilon, d, p)$-lower-regular in G.*

Proof. Let X and Y be arbitrary subsets of U and V with $|X| \geq 2\varepsilon|U|, |Y| \geq 2\varepsilon|V|$. We need to show that $e_G(X, Y) \geq (d - 2\varepsilon)p|X||Y|$. Since (U', V') is (ε, d, p)-lower-regular in G, and since $2\varepsilon - \frac{1}{10}\varepsilon^3 > \varepsilon$, we have

$$e_G(X \cap U', Y \cap V') \geq (d-\varepsilon)p|X \cap U'||Y \cap V'|$$
$$\geq (d-\varepsilon)p\left(|X| - \tfrac{1}{10}\varepsilon^3|U'|\right)\left(|Y| - \tfrac{1}{10}\varepsilon^3|V'|\right) \geq (d-\varepsilon)(1 - \tfrac{1}{10}\varepsilon^2)^2 p|X||Y|$$
$$\geq (d-2\varepsilon)p|X||Y|,$$

as desired. $\square$

Finally, we need a statement about the edge distribution in $G(n, p)$.

Lemma 2.11. *For any $p \geq n^{-1/2}$, a.a.s. $\Gamma = G(n, p)$ has the following property.*

$(\star)$ *For any disjoint $W, Z \subseteq V(\Gamma)$ with $|W| \leq |Z|$ and $|W| \geq 48p^{-1}\log(en/|Z|)$, we have $e_\Gamma(W, Z) = \left(1 \pm \frac{1}{2}\right)p|W||Z|$, and for any $Z \subseteq V(\Gamma)$ we have $e_\Gamma(Z) \leq 4p|Z|^2 + 2|Z|\log n$.*

Proof. Indeed, using Theorem 2.1 and a union bound over the choices of W and Z, we sum over all possible cardinalities of W and Z to obtain that the failure probability of the first assertion is at most

$$\sum_{|W|,|Z|} \binom{n}{|W|}\binom{n}{|Z|} \cdot 2\exp\left(-\frac{p|W||Z|}{12}\right)$$

$$\leq \sum_{|W|,|Z|} 2 \cdot \left(\tfrac{en}{|Z|}\right)^{2|Z|} \exp\left(-4|Z|\log\tfrac{en}{|Z|}\right)$$

$$\leq 2n^2 \exp\left(-n^{-1/2}\right),$$

where the first inequality uses $|W|,|Z| \geq 48p^{-1}\log(en/|Z|)$, and the second our assumption $p \geq n^{-1/2}$. This tends to zero as n tends to infinity, as desired. Similarly, since $4p|Z|^2$ is greater than seven times the expected number of edges in Z, the failure probability of the second assertion is at most

$$\sum_{|Z|} n^{|Z|} \cdot e^{-2|Z|\log n} = \sum_{|Z|} n^{-|Z|},$$

which tends to zero as n tends to infinity. $\qquad\square$

We now prove our one-sided regularity inheritance lemma, Lemma 1.26.

Proof of Lemma 1.26. Given $0 < \varepsilon', d$, we set $\delta = \varepsilon'/2$. We obtain ε_2 and C_1 from Theorem 2.7 for input d, $\beta = e^{100}$ and $\delta^2/10$. Now we obtain ε_1 and C_2 from Lemma 2.9 for input d, ε_2 and δ. We set $\varepsilon_0 = \frac{1}{2}\varepsilon_1$. We set $C' = 1000C_1C_2(\varepsilon')^{-3}$ and $C = 1000(\varepsilon')^{-3}C'$. We may suppose that $p \geq n^{-1/2}$, since otherwise the lemma holds vacuously as it requires $|X| > n$.

We now show that, provided Γ has the property $(\star)$ of Lemma 2.11, given a counterexample (X,Y) and subgraph G of Γ to the conclusion of Lemma 1.26 we can find $X' \subseteq X$, $Y' \subseteq Y$ and a set Z' such that X', Y' and Z' are disjoint, such that (X',Y') is (ε_1,d,p)-regular, and such that each $z \in Z'$ has $\deg_\Gamma(z;X')$ between $\frac{1}{2}p|X'|$ and $\frac{3}{2}p|X'|$, yet the pair $\left(N_\Gamma(z;X'),Y'\right)$ is not (δ,d,p)-regular in G.

To that end, suppose that (X,Y) and a subgraph G of Γ forms a counterexample. Recall that we have $|X| \geq C\max(p^{-2},p^{-1}\log n)$ and $|Y| \geq Cp^{-1}\log(en/|X|)$. Let Z be a set of $C'p^{-1}\log(en/|X|)$ vertices witnessing this counterexample, i.e. a set of vertices whose neighbourhoods in X do not form (ε',d,p)-regular pairs with Y in G. By $(\star)$, we have $e_\Gamma(Z) \leq 4p|Z|^2 + 2|Z|\log n$, so in particular for at least half of the vertices $z \in Z$ we have $\deg_\Gamma(z;Z) \leq 16p|Z| + 8\log n \leq 20C'\log n$. Now let $X' = X \setminus Z$, and $Y' = Y \setminus Z$. By choice of C', we have $|X'| \geq \left(1 - \frac{1}{100}(\varepsilon')^3\right)|X|$, and $|Y'| \geq \left(1 - \frac{1}{100}(\varepsilon')^3\right)|Y|$. Furthermore, by $(\star)$, there are at most $100p^{-1}\log(en/|X'|)$ vertices $z \in Z$ such that $\deg_\Gamma(z;X') \neq \left(1 \pm \frac{1}{2}\right)p|X'|$. Let Z' consist of $\frac{1}{4}C'p^{-1}\log(en/|X|)$ vertices in Z with at most $20C'\log n$ neighbours in Z and $\left(1 \pm \frac{1}{2}\right)p|X'|$ neighbours in X', which we can do by choice of C'. We claim that (X',Y',Z') form our desired smaller counterexample. Indeed, by Lemma 1.15, (X',Y') is (ε_1,d,p)-regular in G. Furthermore, by Lemma 2.10 and choice of C, if $\left(N_\Gamma(z;X'),Y'\right)$ were (δ,d,p)-regular then $\left(N_\Gamma(z;X),Y\right)$ would be (ε',d,p)-regular, in contradiction to the definition of Z. Thus $\left(N_\Gamma(z;X'),Y'\right)$ is not (δ,d,p)-regular for any $z \in Z'$.

We now use Lemma 2.9 to obtain Y^*. We apply Lemma 2.9 to the regular pair (Y',X'), with input d, ε_2 and δ, and with $i = 0$ and $m = C'p^{-1}\log(en/|X|)$. We let $\ell = |Z'|$ and the sets $(X_j,Y_j)_{j\in[\ell]}$ be the pairs $(Y',N_\Gamma(z;X'))$. The result is that there exist choices of $Y^* \subseteq Y'$ with $|Y^*| = m$ such that (X',Y^*) is

(ε_2, d, p)-regular in G, yet for each $z \in Z'$ the pair $\left(N_\Gamma(z; X'), Y^*\right)$ is not $(\delta^2/10, d, p)$-regular in G. We fix one such Y^*.

We have now shown that if for some Γ we have $(\star)$ but the conclusion of Lemma 1.26 fails, then there exist (X', Y^*, Z') and a graph G on (X', Y^*) with the properties above. Our aim now is to show that this *bad object* is unlikely to exist in $G(n, p)$. To that end, fix any $X' \subseteq [n]$ of size at least $\frac{1}{2}C \max(p^{-2}, p^{-1}\log n)$. Fix now disjoint sets $Y^*, Z' \subseteq [n] \setminus X'$ of sizes m and $\frac{1}{4}m$ respectively. We reveal the edges of $G(n, p)$ between X' and Y^*, and let G be any fixed (ε_2, d, p)-regular subgraph of these edges. Next, we reveal the edges of $\Gamma = G(n, p)$ between Z' and X'. We would like to upper bound the probability that we obtain $\deg_\Gamma(z; X') = \left(1 \pm \frac{1}{2}\right) p |X'|$ for each $z \in Z'$, and $\left(N_\Gamma(z; X'), Y^*\right)$ is not $(\delta^2/10, d, p)$-regular for any $z \in Z'$. Since the edges from each $z \in Z'$ to X' are independent, it is enough to estimate the probability that $\left(N_\Gamma(z; X'), Y^*\right)$ fails to be $(\delta^2/10, d, p)$-regular in G, conditioning on $\deg_\Gamma(z; X')$. By Theorem 2.7, this probability is at most $\beta^{\deg_\Gamma(z; X')} \leq \beta^{p|X'|/2}$. We conclude that the probability that a given X', Y^*, Z' and G form a bad object is at most $\beta^{p|X'||Z'|/2}$. We now take the union bound over choices of X' (including the choice of $|X'|$), Y^* and Z' (whose sizes are fixed given $|X'|$), and over the choices of $G \subseteq \Gamma$. Note that because $(\star)$ holds, the number of edges in Γ between X' and Y^* is at most $3p|X'||Y^*|/2$, so the number of choices of $G \subseteq \Gamma$ is at most $2^{3p|X'||Y^*|/2}$. Finally, using the facts that $\frac{1}{2}|X| \leq |X'| \leq |X|$, that $|Y^*| = m$ and $|Z'| = m/4$, and $|X'| \geq m$, we conclude that the probability that $(\star)$ holds but $G(n, p)$ contains a bad object is at most

$$
\sum_{|X'|} \binom{n}{|X'|}\binom{n}{|Y^*|}\binom{n}{|Z'|} 2^{3p|X'||Y^*|/2} \beta^{p|X'||Z'|/2}
$$

$$
\leq \sum_{|X'|} \left(\tfrac{en}{|X'|}\right)^{6|X'|} 2^{3p|X'|m/2} \beta^{p|X'|m/8}
$$

$$
\leq \sum_{|X'|} \exp\left(12|X'|\log\tfrac{en}{|X'|} + \tfrac{3}{2}C'|X'|\log\tfrac{en}{|X'|} - 12.5C'|X'|\log\tfrac{en}{|X'|}\right)
$$

$$
\leq \sum_{|X'|} e^{-|X'|}
$$

where the final sum tends to zero as n tends to infinity since $|X'| \geq \frac{1}{2}Cp^{-1}\log n$, which tends to infinity with n. Since by Claim 2.11, $G(n, p)$ has $(\star)$ a.a.s., we conclude that a.a.s. $\Gamma = G(n, p)$ satisfies the conclusion of Lemma 1.26. □

Our proof of Lemma 1.27 is closely related to the above proof. Again, we begin by assuming that a counterexample (X, Y) and G exists, with $|X| \leq |Y|$, and deducing the existence of disjoint sets X', Y' and Z'. This time we will obtain $|X'| = |Y'|$ to be only slightly smaller than $|X|$. Again, we will have that each vertex of Z' has roughly the expected size neighbourhood in each of X' and Y' but these neighbourhoods fail to induce regular pairs in G. Now, however, we have to work harder to show that this structure is unlikely to exist in $G(n, p)$. In order to make the union bound over choices of G work, we have to apply Lemma 2.9 twice to reduce X' to X^* and Y' to Y^*, and we have to choose $|X^*| = |Y^*| = m$ of comparable size to Z'. Unfortunately now, if $p^{-1}\log(en/|X|)$ is too small the obvious union bound over choices of X^*, Y^* and Z' fails. Thus our *bad object* is not one triple (X^*, Y^*, Z') together with G, but that (X', Y', Z') contains a collection of $\frac{1}{4}\binom{|X'|}{m}\binom{|Y'|}{m}$ such triples (X^*, Y^*, Z'), which we obtain using the full strength of

Lemma 2.9. We will show that this bad object is sufficiently unlikely to exist to make the union bound over choices of X', Y' and Z' succeed.

Proof of Lemma 1.27. Given $0 < \varepsilon', d$, we set $\delta_1 = \frac{1}{2}\varepsilon'$ and for each $2 \leq j \leq 5$ we set $\delta_j = \delta_{j-1}^2/10$. We obtain ε_5 from Corollary 2.8 for input d, $\beta = e^{-100}$ and δ_5. For each $j = 4, \ldots, 1$ we obtain ε_j and C_{j+1} from Lemma 2.9 for input d, ε_{j+1}, δ_j and $i = 1$. We set $\varepsilon_0 = \frac{1}{2}\varepsilon_1$. Finally, we let $C' = 1000C_1C_2C_3C_4C_5(\varepsilon')^{-3}$, and $C = 1000(\varepsilon')^{-3}C'$. Given $p > 0$ and $|X|$, let

$$m = C' \max\left(p^{-2}, p^{-1}\log(en/|X|)\right).$$

We will always apply Lemma 2.9 and Corollary 2.8 with input pairs of sets of sizes at least m, and since $m \geq C'p^{-2}$, by choice of C' this is indeed large enough for these two lemmas. This is the only place where we need $m \geq C'p^{-2}$ in the proof; our probabilistic calculations will always need $m \geq C'p^{-1}\log(en/|X|)$.

Again, we may suppose that $p \geq n^{-1/2}$ since otherwise the lemma statement holds vacuously. We may also assume $|Y| \geq |X|$. As before, we assume that Γ satisfies $(\star)$ but that (X, Y) and the subgraph G of Γ is a counterexample to the conclusion of Lemma 1.27. Our aim is to find disjoint sets X', Y' and Z' which witness failure of regularity inheritance. As before, we let Z consist of $C' \max\left(p^{-2}, p^{-1}\log(en/|X|)\right)$ vertices z such that $\left(N_\Gamma(z;X); N_\Gamma(z;Y)\right)$ fails to be (ε', d, p)-regular in G. We let $X'' = X \setminus Z$, and $Y'' = Y \setminus Z$. As before, using $(\star)$ at least $|Z|/2$ vertices $z \in Z$ have $\deg_\Gamma(z;Z) \leq 20C' \max(p^{-1}, \log n)$, and at most $200p^{-1}\log(en/|X''|)$ vertices of Z either fail $\deg_\Gamma(z;X'') = \left(1 \pm \frac{1}{2}\right)p|X''|$ or $\deg_\Gamma(z;Y'') = \left(1 \pm \frac{1}{2}\right)p|Y''|$. Again, by choice of C we can find a set Z' of $\frac{1}{4}C' \max\left(p^{-2}, p^{-1}\log(en/|X|)\right) = \frac{1}{4}m$ vertices $z \in Z$ which have $\deg_\Gamma(z;Z) \leq 20C' \max(p^{-1}, \log n)$, and $\deg_\Gamma(z;X'') = \left(1 \pm \frac{1}{2}\right)p|X''|$, and $\deg_\Gamma(z;Y'') = \left(1 \pm \frac{1}{2}\right)p|Y''|$. Again, by Lemma 1.15 the pair (X'', Y'') is (ε_1, d, p)-regular in G, and again by Lemma 2.10 for each $z \in Z'$ the pair $\left(N_\Gamma(z;X''), N_\Gamma(z;Y'')\right)$ is not (δ_1, d, p)-regular in G.

We now choose k to be the largest integer such that $km \leq |X''|, |Y''|$. By choice of C we have $km \geq |X|/2$. We apply Lemma 2.9 twice. First we apply it with input d, ε_2, δ_1 and $i = 1$ to (X'', Y'') to obtain $X' \subseteq X''$ with $|X'| = km$. Second, we apply it with input d, ε_3, δ_2 and $i = 1$ to (X', Y'') to obtain $Y' \subseteq Y''$ with $|Y'| = km$. By construction, the pair (X', Y') is (ε_3, d, p)-regular in G, for each $z \in Z'$ we have $\deg_\Gamma(z;X'), \deg_\Gamma(z;Y') = \left(1 \pm \frac{1}{2}\right)km$, and for each $z \in Z'$ the pair $\left(N_\Gamma(z;X'), N_\Gamma(z;Y')\right)$ is not (δ_3, d, p)-regular in G.

We now show that this triple (X', Y', Z') contains the claimed bad object. Again, we need to apply Lemma 2.9. This time, we first apply it with input d, ε_4, δ_3 and $i = 1$ to (X', Y') to obtain $\frac{1}{2}\binom{km}{m}$ choices of X^*, each of size m. For each such choice of X^*, we apply it with input d, ε_5, δ_4 and $i = 1$ to (X^*, Y') to obtain $\frac{1}{2}\binom{km}{m}$ choices of Y^*, each of size m. In total, we find $\frac{1}{4}\binom{km}{m}^2$ choices of (X^*, Y^*) which is an (ε_5, d, p)-lower-regular pair in G with parts of size m, such that for each $z \in Z'$ we have $\deg_\Gamma(z;X^*), \deg_\Gamma(z;Y^*) = \left(1 \pm \frac{1}{2}\right)^2 pm$ and $\left(N_\Gamma(z;X^*), N_\Gamma(z;Y^*)\right)$ is not (δ_5, d, p)-lower-regular in G. We call such a pair (X^*, Y^*) a *bad pair* with respect to Z', and the triple (X', Y', Z') containing $\frac{1}{4}\binom{km}{m}^2$ bad pairs the *bad object*.

We now show that a.a.s. $G(n, p)$ does not contain any such bad object. To that end, with k given we fix disjoint sets X', Y', Z' of size $|Z'| = m$ and $|X'| = |Y'| = km$. Suppose that for some graph G this triple forms a bad object. Picking a uniform random partition $X' = X_1^* \cup \cdots \cup X_k^*$ and $Y' = Y_1^* \cup \cdots \cup Y_k^*$, in expectation at least $k/4$ of the pairs (X_j^*, Y_j^*) are bad pairs. In particular, letting T be the number of pairs

of partitions of X' and Y' into sets of size m, we see that at least $T/8$ pairs of partitions contain $k/8$ or more bad pairs.

On the other hand, we can estimate the probability that there exists a graph $G \subseteq G(n,p)$ on (X',Y',Z') such that the uniform random partition contains at least $k/8$ bad pairs in $G(n,p)$. We expose the edges of $\Gamma = G(n,p)$ between each X_j^* and Y_j^*, and fix any (ε_5, d, p)-lower-regular subgraphs to obtain G. Now we expose the edges from Z' to each of X' and Y'. When exposing the edges for a $z \in Z'$ we first reveal its degrees. Note that if the degree of z to some X_j^* or Y_j^* is not $(1 \pm \frac{1}{2})^2 pm$, then by definition (X_j^*, Y_j^*) is not a bad pair with respect to Z'. Conditioned on the event that all degrees are within $(1 \pm \frac{1}{2})^2 pm$ Corollary 2.8 implies that the probability that for any given $z \in Z'$ and $j \in [k]$ the pair $\big(N_\Gamma(z; X_j^*), N_\Gamma(z; Y_j^*)\big)$ is not (δ_5, d, p)-regular is at most $\beta^{pm/4}$. Since these events are independent for different $z \in Z'$, the probability that any given (X_j^*, Y_j^*) forms a bad pair with respect to Z' is at most $\beta^{pm^2/16}$. If Γ has property $(\star)$, the probability that there exists some graph $G \subseteq \Gamma$ such that (X_j^*, Y_j^*) forms a bad pair in G with respect to Z' is thus at most $2^{3pm^2/2}\beta^{pm^2/16}$, which by choice of β is at most $\beta^{pm^2/32}$.

Thus the probability that for at least $k/8$ of these pairs there is a graph G in which they are bad with respect to Z' is

$$\sum_{\ell=k/8}^{k} \binom{k}{\ell} \beta^{\ell pm^2/16} \big(1 - \beta^{pm^2/32}\big)^{k-\ell}$$

$$\leq k \cdot \binom{k}{k/8} \beta^{kpm^2/1000} \leq 2^k \beta^{pkm^2/1000}$$

where the first inequality is by choice of β. By linearity of expectation, the expected number of pairs of partitions of X' and Y' in $G(n,p)$ for which there is a graph G such that at least $k/8$ pairs are bad with respect to Z' is at most $2^k \beta^{pkm^2/1000} \cdot T$. Applying Markov's inequality, we see that the probability that (X', Y', Z') can be a bad object with respect to any graph G is at most $8 \cdot 2^k \beta^{pkm^2/2000}$. Finally, taking the union bound over choices of k, X', Y' and Z' we see that the probability of $G(n,p)$ containing a bad object is at most

$$\sum_{k=1}^{n/m} \binom{n}{km}^2 \binom{n}{m} \cdot 8 \cdot 2^k \beta^{pkm^2/2000} \leq 8 \sum_{k=1}^{n/m} \big(\tfrac{n}{km}\big)^{3km} 2^k \beta^{pkm^2/2000}$$

$$\leq 8 \sum_{|X|=\log n}^{n} \exp\big(3|X|\log(en/|X|) + |X|/m + (\log \beta) \cdot C'|X|\log(en/|X|)/2000\big)$$

$$\leq \sum_{|X|=\log n}^{n} \exp\big(-|X|\log(en/|X|)\big)$$

where the third inequality is since $|X|/2 \leq km \leq |X|$ and $|X| \geq \log n$, and the fourth by choice of β and C'. This tends to zero as n tends to infinity, which together with Lemma 2.11 shows that a.a.s. $G(n,p)$ does not contain any bad object. This completes the proof of Lemma 1.27. $\qquad\square$

We conclude this section by sketching constructions which show that the number of vertices failing to inherit lower-regularity in each of Lemmas 1.26 and 1.27 are sharp for all $n^{-1/2} \leq p \leq \frac{1}{2}$ and all $C\max(p^{-2}, p^{-1}\log n) \leq |X| \leq n/4$.

First, for Lemma 1.26, given $C \max(p^{-2}, p^{-1} \log n) \le m \le n/100$ we fix a set Z of $\left\lfloor \frac{1}{2} p^{-1} \log(en/m) \right\rfloor$ vertices. If $|Z| = 0$, then trivially there is a set X of m vertices with no neighbours in Z. If $|Z| \ge 1$, then the probability that a vertex outside Z has no neighbours in Z is

$$(1-p)^{|Z|} \ge e^{-2p \cdot p^{-1} \log(en/m)/4} = \left(\frac{m}{en}\right)^{1/2}$$

and so the expected number of vertices outside Z with no neighbours in Z is at least $\frac{99n}{100}(m/en)^{1/2} \ge 2m$. By Theorem 2.1, since m tends to infinity with n, a.a.s. there is a set X of m vertices with no neighbours in Z. Furthermore, for each $\varepsilon > 0$ a.a.s. the pair $(X, [n] \setminus X)$ is $(\varepsilon, \frac{1}{2}, p)$-lower-regular, but trivially for each vertex $z \in Z$ the neighbourhood of z in X does not form a regular pair with $[n] \setminus X$.

For Lemma 1.27, the same construction shows we need $|Z| \ge cp^{-1} \log(en/|X|)$. Finally, for each $\varepsilon > 0$ there exists $p_\varepsilon > 0$ such that for $0 < p < p_\varepsilon$, Huang, Lee and Sudakov [45, Proposition 6.3] construct a subgraph G of Γ with minimum degree $(1-\varepsilon)pn$ in which $\Omega(p^{-2})$ vertices are not in triangles. This construction in particular shows that $\Omega(p^{-2})$ vertices can fail to inherit regularity in Lemma 1.27.

2.2 Deterministic properties of the ambient graph

In this section we introduce the deterministic properties which we require of our ambient graphs Γ. We then also prove that random and bijumbled graphs have (a subset of) these properties, in Lemma 2.17 and Lemma 2.18, respectively. These are exactly the properties which we shall use in our blow-up lemma proofs.

Our first deterministic property asserts that most vertices of Γ have close to the expected degree into any given reasonably large subset of vertices. This property will be used for both random and bijumbled graphs.

Definition 2.12 (Neighbourhood size property $NS(\varepsilon, T, \Delta)$). Given $\varepsilon > 0$ and integers T and Δ, we say that the graph Γ has property $NS(\varepsilon, T, \Delta)$ if the following is true for some p such that Γ has density $(1 \pm \varepsilon)p$. For any set W of size at least $\varepsilon p^{\Delta-1} v(\Gamma)/T^2$, there are at most $\varepsilon p^{\Delta-1} v(\Gamma)/T^2$ vertices v outside W such that $\deg_\Gamma(v; W) \ne (1 \pm \varepsilon)p|W|$.

Our next property concerns one-sided and two-sided regularity inheritance. Based on the regularity inheritance lemmas established in Section 2.1.3 we shall also establish this property for random as well as bijumbled graphs. Recall our convention though, that in random graphs regular pairs are lower-regular pairs and in bijumbled graphs they are fully-regular pairs.

Definition 2.13 (Regularity inheritance property $RI(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, T, \Delta)$).
Given $0 < \varepsilon < \varepsilon' < d$ and integers T, Δ, we say that the graph Γ has property $RI(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, T, \Delta)$ if we have $\varepsilon_{a,b} \in [\varepsilon, \varepsilon']$ for each $0 \le a, b \le \Delta - 1$ such that the following holds for some p such that Γ has density $(1 \pm \varepsilon)p$.

If $0 \le a \le \Delta - 2$ and $0 \le b \le \Delta - 1$, if $G \subseteq \Gamma$ and if $X, Y \subseteq V(\Gamma)$ are disjoint sets with

$$|X| \ge \varepsilon' p^{\Delta-2} \cdot \frac{v(\Gamma)}{T^2} \qquad \text{and} \qquad |Y| \ge \varepsilon' p^{\Delta-1} \cdot \frac{v(\Gamma)}{T^2}$$

and if (X,Y) is $(\varepsilon_{a,b},d,p)$-regular pair in G, then $\big(N_\Gamma(v;X),Y\big)$ is $(\varepsilon_{a+1,b},d,p)$-regular in G for all but at most $\varepsilon p^{\Delta-1}v(\Gamma)/T^2$ vertices $v \in V(\Gamma) \setminus (X \cup Y)$. Furthermore, if additionally $b \leq \Delta - 2$ and

$$|Y| \geq \varepsilon' p^{\Delta-2} \cdot \frac{v(\Gamma)}{T^2}\,,$$

then the pair $\big(N_\Gamma(v;X),N_\Gamma(v;Y)\big)$ is $(\varepsilon_{a+1,b+1},d,p)$-regular in G for all but at most $\varepsilon p^{\Delta-2}v(\Gamma)/T^2$ vertices $v \in V(\Gamma) \setminus (X \cup Y)$.

The next property concerns the count of certain stars in Γ, and we shall call it congestion property (following [54], where a very similar property was used). We only establish this property for random graphs. For bijumbled graphs it does not hold (with any reasonable choice of parameters). Given a graph Γ, a set $U \subseteq V(\Gamma)$ and a collection $\mathcal{F}$ of pairwise disjoint ℓ-sets in $V(\Gamma)$, we define the *congestion graph* $\mathrm{CG}(\Gamma,U,\mathcal{F})$ to be the bipartite graph with vertex sets U and $\mathcal{F}$ with uF an edge of $\mathrm{CG}(\Gamma,U,\mathcal{F})$ if $u \in U$ is a common neighbour in Γ of the vertices in $F \in \mathcal{F}$.

Definition 2.14 (Congestion property $\mathrm{CON}(\rho,T,\Delta)$)**.** Given $\rho > 0$ and integers T and Δ, we say that Γ has property $\mathrm{CON}(\rho,T,\Delta)$ if the following statement is true for some p such that Γ has the density $(1 \pm \rho)p$. For each $1 \leq \ell \leq \Delta$, each $U \subseteq V(\Gamma)$ and each collection $\mathcal{F}$ of pairwise disjoint ℓ-sets in $V(\Gamma) \setminus U$, if we have $|U| \leq |\mathcal{F}| \leq \rho v(\Gamma)$, then

$$e\big(\mathrm{CG}(\Gamma,U,\mathcal{F})\big) \leq 7p^\ell |U||\mathcal{F}| + \rho p^\ell \cdot \frac{v(\Gamma)}{T}|\mathcal{F}|\,.$$

The congestion condition will help us to verify Hall's condition (on some linearly sized set) in order to embed many vertices at a time. However, in the proof of our blow-up lemma for degenerate graphs, Lemma 1.23, we cannot use this strategy. Instead we will embed the vertices one by one in the given order. For this purpose we need to have a 'local' version of the congestion property, which we will want to apply to sets U in the common Γ-neighbourhood of already embedded vertices. We remark that for such small sets U the bound on the number of edges in the congestion graph given by the congestion property becomes trivial. The following local congestion property is designed to give a useful bound in this case. Again, we shall only establish this property for random graphs.

Definition 2.15 (Local congestion property $\mathrm{LCON}(\varepsilon,T,\Delta)$)**.** Given $\varepsilon > 0$ and integers T and Δ, we say that Γ has property $\mathrm{LCON}(\varepsilon,T,\Delta)$ if the following statement is true for some p such that Γ has density $(1 \pm \varepsilon)p$. For each $i,\ell \geq 1$ with $i + \ell \leq \Delta$, each $U \subseteq V(\Gamma)$ of size at least $\varepsilon p^i v(\Gamma)/T^2$ and each collection $\mathcal{F}$ of pairwise disjoint ℓ-sets in $V(\Gamma) \setminus U$ we have

$$e\big(\mathrm{CG}(\Gamma,U,\mathcal{F})\big) \leq 7p^\ell |U| \max\big(\varepsilon|U|,|\mathcal{F}|\big)\,.$$

Since for bijumbled graphs we cannot use the congestion property or local congestion property we introduce the following final deterministic property which we shall use there instead. This property is a strengthening of the neighbourhood size property, which also takes smaller sets W in which we measure neighbourhoods into consideration, and further requires the number of exceptional vertices outside W to be smaller. We call it 'lopsided' to distinguish it from property $\mathrm{NS}(\varepsilon,T,\Delta)$, in which the maximum number of failing vertices and the minimum size of the set into which neighbourhoods are taken are the same.

Definition 2.16 (Lopsided neighbourhood size property $\mathrm{LNS}(\varepsilon,T,\Delta)$)**.** Given $\varepsilon > 0$ and integers T and Δ, we say that Γ has property $\mathrm{LNS}(\varepsilon,T,\Delta)$ if the following is true for some p such that Γ has density $(1 \pm \varepsilon)p$. For each $0 \leq j \leq \Delta - 1$ and any set W of size at least $\varepsilon p^{\Delta+j} v(\Gamma)/T^2$, there are at most $\varepsilon p^{2\Delta-j-1} v(\Gamma)/T^2$ vertices v outside W such that $\deg_\Gamma(v;W) \neq (1 \pm \varepsilon)p|W|$.

2.2.1 Random graphs

In this subsection we will show that $G(n,p)$ has the neighbourhood size property, regularity inheritance property, congestion property and local congestion property by proving the following lemma

Lemma 2.17 (Deterministic properties of $G(n,p)$)**.** *For every $\Delta \geq 2$ and $d, \varepsilon' > 0$ there exist $\varepsilon > 0$ and $\varepsilon_{a,b} > 0$ for each $0 \leq a,b \leq \Delta - 1$ such that for every T and $\rho > 0$ there exists $C > 0$ such that if $p \geq C(\log n/n)^{1/\Delta}$ then $G(n,p)$ a.a.s. has*

(a) $\mathrm{NS}(\varepsilon,T,\Delta)$,

(b) $\mathrm{RI}(\varepsilon,(\varepsilon_{a,b}),\varepsilon',d,T,\Delta)$,

(c) $\mathrm{CON}(\rho,T,\Delta)$,

(d) $\mathrm{LCON}(\varepsilon,T,\Delta)$.

The proof of part (a) is standard. Part (b) follows from the regularity inheritance lemmas, Lemmas 1.26 and 1.27. The proof of part (c) follows arguments from [54] (but is slightly different), and part (d) is proved similarly. For completeness we give the details.

Proof of Lemma 2.17. Given $\Delta \geq 1$ and $d, \varepsilon' > 0$, we assume without loss of generality that $(1 - \varepsilon')^\Delta > 1/2$. We choose $\varepsilon_{a,b}$ for each $0 \leq a \leq \Delta - 1$ and $0 \leq b \leq \Delta - 1$ as follows. We set $\varepsilon_{\Delta-1,\Delta-1} = \varepsilon'/2$ and we define the other $\varepsilon_{a,b}$ inductively. For each a and b, we require that $\varepsilon_{a,b}$ is smaller than the ε_0 returned by Lemma 1.26 with input $\varepsilon_{a+1,b}/2$ and d (provided $a < \Delta - 1$), and that returned with input $\varepsilon_{a,b+1}/2$ and d (provided $b < \Delta - 1$), and also than the ε_0 returned by Lemma 1.27 with input $\varepsilon_{a+1,b+1}/2$ and d (provided $a < \Delta - 1, b < \Delta - 1$). Let ε be the minimum of the $\varepsilon_{a,b}$. Note that we then have $\varepsilon_{a,b} = \varepsilon_{b,a}$ for each a,b. Given T, let C_1 be the maximum of all of the constants C returned by the above applications of Lemmas 1.26 and 1.27. Further, given ρ set

$$C = \frac{100T^2 C_1 \Delta}{\varepsilon^4 \rho}.$$

Now let $p \geq C(\log n/n)^{1/\Delta}$. Let $\Gamma = G(n,p)$. By the Chernoff bound in Theorem 2.1, since $p > 1/n$, a.a.s. Γ has density $(1 \pm \min(\varepsilon,\rho))p$. In the following we condition on this occurring.

Proof of (a): Using the Chernoff bound again, if X and Y are any disjoint sets of size at least $6\varepsilon^{-2}p^{-1}\log n$, then the probability that $e(X,Y) \neq (1 \pm \varepsilon)p|X||Y|$ is at most $2\exp(-\varepsilon^2 p|X||Y|/3)$. It

follows that the probability that there exist such X and Y is at most

$$\sum_{|X|,|Y|\geq 6\varepsilon^{-2}p^{-1}\log n} n^{|X|+|Y|} \cdot 2\exp\Big(-\frac{\varepsilon^2 p|X||Y|}{3}\Big)$$

$$\leq \sum_{|X|,|Y|} 2\cdot 2^{2\max(|X|,|Y|)\log n}\exp\big(-2\max(|X|,|Y|)\log n\big)$$

$$\leq 2n^2\big(2^2\exp(-2)\big)^{\log^2 n},$$

where the first inequality uses $\min(|X|,|Y|) \geq 6\varepsilon^{-2}p^{-1}\log n$. It follows that a.a.s. any such pair of sets in Γ has density $(1\pm\varepsilon)p$. We condition on this in the following.

Now suppose that X is any set of size at least $\varepsilon p^{\Delta-1}n/T^2$ in $V(\Gamma)$, and Y is the set of vertices outside X with fewer than $(1-\varepsilon)p|X|$ neighbours in X. Since the density of (X,Y) is less than $(1-\varepsilon)p$, and since by choice of C we have $\varepsilon p^{\Delta-1}n/T^2 > 6\varepsilon^{-2}p^{-1}\log n$, we conclude that $|Y| < 6\varepsilon^{-2}p^{-1}\log n < \varepsilon p^{\Delta-1}n/(2T^2)$ (again by choice of C). The same argument bounds the number of vertices outside X with more than $(1+\varepsilon)p|X|$ neighbours in X, so Γ has $\mathrm{NS}(\varepsilon,T,\Delta)$ as desired.

Proof of (b): In this proof we shall apply the regularity inheritance lemmas, Lemmas 1.26 and 1.27, less than $3\Delta^2$ times. Hence we can assume that the less than $3\Delta^2$ properties asserted to hold a.a.s. for $G(n,p)$ by these applications of the lemmas hold simultaneously a.a.s for $G(n,p)$. Note that whenever we apply these lemmas, we have $|X| > e$ and hence $\log(en/|X|) \leq \log n$. We will not need to use the sharper bound taking into account the size of X given in Lemmas 1.26 and 1.27.

Given $0 \leq a \leq \Delta-2$ and $0 \leq b \leq \Delta-1$, suppose that (X,Y) is an $(\varepsilon_{a,b},d,p)$-lower-regular pair in a subgraph G of Γ with $|X| \geq \varepsilon' p^{\Delta-2}n/T^2$ and $|Y| \geq \varepsilon' p^{\Delta-1}n/T^2$. Let Z be the set of vertices outside $X \cup Y$ such that $(N_\Gamma(z;X),Y)$ is not $(\varepsilon_{a+1,b},d,p)$-lower-regular in G. By choice of p, we have $\varepsilon' p^{\Delta-2}n/T^2 \geq C_1 p^{-2}\log n$ and $\varepsilon' p^{\Delta-1}n/T^2 \geq C_1 p^{-1}\log n$. Therefore we can apply Lemma 1.26 to conclude that

$$|Z| \leq C_1 p^{-1}\log n \leq \tfrac{\varepsilon}{2T^2}Cp^{-1}\log n < \varepsilon p^{\Delta-1}n/T^2,$$

as desired.

Similarly, given $0 \leq a \leq \Delta-2$ and $0 \leq b \leq \Delta-2$, suppose that (X,Y) is an $(\varepsilon_{a,b},d,p)$-lower-regular pair in a subgraph G of Γ with $|X| \geq \varepsilon' p^{\Delta-2}n/T^2$ and $|Y| \geq \varepsilon' p^{\Delta-2}n/T^2$. Let Z be the set of vertices outside $X \cup Y$ such that $(N_\Gamma(z;X),N_\Gamma(z;Y))$ is not $(\varepsilon_{a+1,b+1},d,p)$-lower-regular in G. By choice of p, we have $\varepsilon' p^{\Delta-2}n/T^2 > C_1 p^{-2}\log n$. Therefore we can apply Lemma 1.27 to conclude that

$$|Z| \leq C_1\max(p^{-2},p^{-1}\log n) \leq \tfrac{\varepsilon}{2T^2}C\max(p^{-2},p^{-1}\log n) < \varepsilon p^{\Delta-2}n/T^2,$$

again as desired.

Proof of (c): Given $1 \leq i \leq \Delta$, a set $U \subseteq V(\Gamma)$, and a family $\mathcal{F}$ of pairwise disjoint ℓ-sets in $V(\Gamma) \setminus U$ with $|U| \leq |\mathcal{F}| \leq \rho|V(\Gamma)|$, the graph $\mathrm{CG}(\Gamma,U,\mathcal{F})$ is a random bipartite graph with edge probability p^ℓ and with parts U and $\mathcal{F}$. So the expected number of edges of $\mathrm{CG}(\Gamma,U,\mathcal{F})$ is $p^\ell|U||\mathcal{F}|$. By the Chernoff bound in Theorem 2.1 the probability that

$$\big|e\big(\mathrm{CG}(\Gamma,U,\mathcal{F})\big)\big| > 7p^\ell|U||\mathcal{F}| + \rho p^\ell n|\mathcal{F}|/T$$

is at most

$$\exp\left(-\rho p^\ell n|\mathcal{F}|/T\right).$$

If $|\mathcal{F}| = m \leq \rho n$, then since $|U| \leq |\mathcal{F}|$ we have $|U| \leq m$. Taking a union bound over ℓ, over m, over the at most n^{m+1} choices for U and its size, and over the at most $n^{\Delta m}$ choices for $\mathcal{F}$, we see that the probability of failure of $\mathrm{CON}(\rho, T, \Delta)$ is at most

$$\sum_{\ell=1}^{\Delta} \sum_{m=1}^{\rho n} n^{m+1} n^{\Delta m} \exp\left(-\rho p^\ell nm/T\right) \leq \Delta \sum_{m=1}^{\rho n} \exp\left(2\Delta m \log n - \rho p^\Delta nm/T\right)$$
$$\leq \Delta \sum_{m=1}^{\rho n} \exp\left(-(\Delta+3)m \log n\right) < \Delta n \exp\left(-(\Delta+2)\log n\right) < 1/n,$$

where the second inequality uses $p^\Delta n \geq C \log n$ and the choice of C. We conclude that $G(n,p)$ a.a.s. has $\mathrm{CON}(\rho, T, \Delta)$ as desired.

Proof of (d): Given integers $i, \ell \geq 1$ such that $i + \ell \leq \Delta$, a set $U \subseteq V(\Gamma)$ of size at least $\varepsilon p^i n/T^2$, and a family of pairwise disjoint ℓ-sets $\mathcal{F}$ in $V(\Gamma) \setminus U$, the graph $\mathrm{CG}(\Gamma, U, \mathcal{F})$ is a random bipartite graph with edge probability p^ℓ and parts U and $\mathcal{F}$. The expected number of edges it contains is therefore $p^\ell |U||\mathcal{F}|$.

We separate two cases. If $|\mathcal{F}| \leq \varepsilon|U|$, then we have $\mathbb{E}e\left(\mathrm{CG}(\Gamma, U, \mathcal{F})\right) \leq \varepsilon p^\ell |U|^2$, so by Theorem 2.1 the probability of $e\left(\mathrm{CG}(\Gamma, U, \mathcal{F})\right) \geq 7\varepsilon p^\ell |U|^2$ is at most

$$\exp\left(-6\varepsilon p^\ell |U|^2\right) \leq \exp\left(-6\varepsilon^2 p^{i+\ell} n|U|/T^2\right) \leq \exp(-3\Delta|U|\log n).$$

Taking the union bound over the at most Δ^2 choices of i and ℓ, at most n^2 choices of $|U|$ and $|\mathcal{F}| \leq |U|$, the at most $n^{|U|}$ choices of U and the at most $n^{\Delta|U|}$ choices of $\mathcal{F}$, we see that the probability that there exists such a choice with $e\left(\mathrm{CG}(\Gamma, U, \mathcal{F})\right) \geq 7\varepsilon p^\ell |U|^2$ is at most

$$\Delta^2 n^2 n^{(\Delta+1)|U|} \exp(-3\Delta|U|\log n)$$

which tends to zero as n tends to infinity.

Next we consider the case $|\mathcal{F}| \geq \varepsilon|U|$. Again by Theorem 2.1 the probability of $e\left(\mathrm{CG}(\Gamma, U, \mathcal{F})\right) \geq 7p^\ell |U||\mathcal{F}|$ is at most

$$\exp\left(-6p^\ell |U||\mathcal{F}|\right) \leq \exp\left(-6\varepsilon p^{i+\ell} n|\mathcal{F}|/T^2\right) \leq \exp(-3\varepsilon^{-1}\Delta|\mathcal{F}|\log n).$$

Taking the union bound over the at most Δ^2 choices of i and ℓ, at most n^2 choices of $|U|$ and $|\mathcal{F}|$, the at most $n^{|\mathcal{F}|/\varepsilon}$ choices of U and the at most $n^{\Delta|\mathcal{F}|}$ choices of $\mathcal{F}$, we see that the probability that there exists such a choice with $e\left(\mathrm{CG}(\Gamma, U, \mathcal{F})\right) \geq 7p^\ell |U||\mathcal{F}|$ is at most

$$\Delta^2 n^2 n^{(\Delta+1/\varepsilon)|\mathcal{F}|} \exp(-3\varepsilon^{-1}\Delta|\mathcal{F}|\log n)$$

which tends to zero as n tends to infinity. We conclude that $G(n,p)$ a.a.s. has $\mathrm{LCON}(\varepsilon, T, \Delta)$ as desired. $\quad\square$

2.2.2 Bijumbled graphs

In this subsection we will verify the neighbourhood size property, the regularity inheritance property, and the lopsided neighbourhood size property for bijumbled graphs, as stated in the following lemma.

Lemma 2.18 (Deterministic properties of bijumbled graphs). *For each $\Delta \in \mathbb{N}$ and $d, \varepsilon' > 0$ there exist $\varepsilon_{a,b} > 0$ for $0 \leq a, b \leq \Delta$ and $\varepsilon > 0$ such that for each $T \in \mathbb{N}$ there is $c > 0$ such that if $p > 0$ then any (p, β)-bijumbled graph on n vertices has*

(a) $\mathrm{NS}(\varepsilon, T, \Delta + 1)$ *if $\beta \leq cp^{\Delta+1}n$,*

(b) $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, T, \Delta + 1)$ *if $\beta \leq cp^{\Delta+2}n$,*

(c) $\mathrm{LNS}(\varepsilon, T, \Delta)$ *if $\beta \leq cp^{(3\Delta+1)/2}n$.*

Let us briefly justify why we cannot use the congestion property in bijumbled graphs. Indeed, *blowing up* a bijumbled graph by a factor Δ—that is, replacing vertices with independent Δ-sets and edges with complete bipartite graphs—degrades the bijumbledness only slightly, but in this blow up the congestion condition fails badly: If we choose $\ell = \Delta$ and as Δ-tuples the blow ups of vertices, then the congestion property fails by a factor $p^{\Delta-1}$.

In the proof of Lemma 2.18, Part (b) is a consequence of the regularity inheritance lemmas, Lemmas 2.5 and 2.6. Parts (a) and (c) use the following easy consequences of the definition of bijumbled graphs.

Proposition 2.19. *Let Γ be a (p, β)-bijumbled graph on n vertices and let $W \subseteq V(\Gamma)$.*

(a) *If $\varepsilon \geq 2\beta \frac{1}{pn}$, then $e(\Gamma) = (1 \pm \varepsilon)p\binom{n}{2}$.*

(b) *$\deg_\Gamma(v; W) \neq (1 \pm \varepsilon)p|W|$ for at most $2\beta^2/(\varepsilon^2 p^2 |W|)$ vertices $v \in V(\Gamma) \setminus W$.*

Proof. From the definition of jumbledness (1.1) we obtain

$$2e(\Gamma) = e_\Gamma(V, V) = pn^2 \pm \beta n = pn^2 \pm \varepsilon pn^2/2$$

and (a) follows.

For (b) let U be the set of $v \in V(\Gamma) \setminus W$ with $\deg_\Gamma(v; W) > (1+\varepsilon)p|W|$. We have

$$(1+\varepsilon)p|U||W| < e(U, W) \leq p|U||W| + \beta\sqrt{|U||W|}$$

and hence $|U| < \beta^2/(\varepsilon^2 p^2 |W|)$. Similarly, the number of $v \in V(\Gamma) \setminus W$ with $\deg_\Gamma(v; W) < (1-\varepsilon)p|W|$ is also at most $\beta^2/(\varepsilon^2 p^2 |W|)$ and (b) follows. $\qquad\square$

We can now prove Lemma 2.18.

Proof of Lemma 2.18. Given $\Delta \geq 2$ and $d, \varepsilon' > 0$, we assume $(1-\varepsilon')^\Delta > 1/2$. We choose $\varepsilon_{a,b}$ for each $0 \leq a \leq \Delta$ and $0 \leq b \leq \Delta$ as follows. We set $\varepsilon_{\Delta,\Delta} = \varepsilon'/2$ and we define the other $\varepsilon_{a,b}$ inductively. For each a and b, we require that $\varepsilon_{a,b}$ is smaller than the ε_0 returned by Lemma 2.5 with input $\varepsilon_{a+1,b}/2$ and d (provided $a < \Delta$), and that returned with input $\varepsilon_{a,b+1}/2$ and d (provided $b < \Delta$), and the ε_0 returned by Lemma 2.6 with input $\varepsilon_{a+1,b+1}/2$ and d (provided $a < \Delta, b < \Delta$). Let ε be the minimum of the $\varepsilon_{a,b}$. Note

that we then have $\varepsilon_{a,b} = \varepsilon_{b,a}$ for each a,b. Let c' be the minimum of all the constants c_0 returned by the applications of Lemmas 2.5 and 2.6. Given T, let

$$c = \tfrac{1}{2}\varepsilon^2\varepsilon'c'T^{-2}.$$

Let Γ be a (p,β)-bijumbled graph on n vertices with $\beta \leq cp^{\Delta+1}n$. By Proposition 2.19($a$) the graph Γ has density $(1\pm\varepsilon)p$.

Proof of (a): Let $W \subseteq V(\Gamma)$ be a set with $|W| \geq \varepsilon p^\Delta n/T^2$. By Proposition 2.19($b$) the number of vertices $v \in V(\Gamma)\setminus W$ such that $\deg(v;W) \neq (1\pm\varepsilon)p|W|$ is at most

$$\frac{2\beta^2}{\varepsilon^2p^2\cdot\varepsilon p^\Delta nT^{-2}} \leq \frac{2c^2T^2p^\Delta n}{\varepsilon^3} \leq \varepsilon p^\Delta \frac{n}{T^2},$$

and hence Γ has $\mathrm{NS}(\varepsilon,T,\Delta+1)$.

Proof of (c): Assume $\beta \leq cp^{(3\Delta+1)/2}n$. Let $0 \leq j \leq \Delta$ and $W \subseteq V(\Gamma)$ be a set with $|W| \geq \varepsilon p^{\Delta+j}n/T^2$. By Proposition 2.19($b$) the number of vertices $v \in V(\Gamma)\setminus W$ such that $\deg(v;W) \neq (1\pm\varepsilon)p|W|$ is at most

$$\frac{2\beta^2}{\varepsilon^2p^2\cdot\varepsilon p^{\Delta+j}nT^{-2}} \leq \frac{2c^2T^2p^{2\Delta-j-1}n}{\varepsilon^3} \leq \varepsilon p^{2\Delta-j-1}\frac{n}{T^2},$$

and hence Γ has $\mathrm{LNS}(\varepsilon,T,\Delta)$.

Proof of (b): We shall prove that Γ has $\mathrm{RI}(\varepsilon,(\varepsilon_{a,b}),\varepsilon',d,T,\Delta+1)$ if $\beta \leq cp^{\Delta+2}n$ by contradiction. Given $0 \leq a \leq \Delta-1$ and $0 \leq b \leq \Delta$, suppose that (X,Y) is an $(\varepsilon_{a,b},d,p)$-fully-regular pair in $G \subseteq \Gamma$ with $|X| \geq \varepsilon'p^{\Delta-1}n/T^2$ and $|Y| \geq \varepsilon'p^\Delta n/T^2$. Let Z be the set of vertices $z \in V(\Gamma)\setminus(X\cup Y)$ such that $(N_\Gamma(z;X),Y)$ is not $(\varepsilon_{a+1,b},d,p)$-fully-regular in G. Assume for contradiction that $|Z| \geq \varepsilon p^\Delta n/T^2$. Then we have $\max\big(|X||Y|,|X||Z|\big) \geq \varepsilon\varepsilon'p^{2\Delta-1}n^2/T^4$. Therefore, by choice of c and since $(\log_2\frac{1}{p})^{-1} > p$, we see that

$$\beta \leq cp^{\Delta+2}n \leq c'\varepsilon\varepsilon'p^{\Delta+2}n/T^2 < c'p^2(\log_2\tfrac{1}{p})^{-1/2}\sqrt{\max\big(|X||Y|,|X||Z|\big)},$$

and thus, Γ is bijumbled enough to apply Lemma 2.5, whose statement contradicts the assumption on Z.

Similarly, given $0 \leq a,b \leq \Delta-1$, suppose that (X,Y) is an $(\varepsilon_{a,b},d,p)$-fully-regular pair in $G \subseteq \Gamma$ with $|X| \geq \varepsilon'p^{\Delta-1}n/T^2$ and $|Y| \geq \varepsilon'p^{\Delta-1}n/T^2$. Let Z be the set of vertices $z \in V(\Gamma)\setminus(X\cup Y)$ such that $(N_\Gamma(z;X),N_\Gamma(z;Y))$ is not $(\varepsilon_{a+1,b+1},d,p)$-fully-regular in G. Again, assuming for contradiction that $|Z| \geq \varepsilon p^{\Delta-1}n/T^2$, we have

$$\max\big(|X||Y|,|Y||Z|,|X||Z|\big) \geq \varepsilon\varepsilon'p^{2\Delta-2}n^2/T^4,$$

so that by choice of c, we see that

$$\beta \leq cp^{\Delta+2}n \leq c'\varepsilon'p^{\Delta+2}n/T^2 < c'p^3\sqrt{\max\big(|X||Y|,|Y||Z|,|X||Z|\big)},$$

hence the conclusion of Lemma 2.6 contradicts the assumption on Z. $\qquad\square$

2.3 The general setup

This section constitutes the first step towards the proofs of our blow-up lemmas. More or less, one can think of this section as showing that, for each of our blow-up lemmas, it suffices to prove a corresponding blow-up lemma in which several extra conditions are required. This reduction, as mentioned in the proof overview (Section 1.4) involves refining the partitions $\mathcal{X}^{\mathrm{BL}}$ of H and $\mathcal{V}^{\mathrm{BL}}$ of G given by the user of the blow-up lemma to obtain new partitions $\mathcal{X}$ of H and $\mathcal{V}$ of G, which in turn entails replacing the reduced graphs R_{BL} and R'_{BL} with new graphs R and R', and altering various constants mentioned in the blow-up lemmas. It is convenient to keep using the letter choices seen in our blow-up lemmas for these altered partitions, graphs and constants in the proofs to come, so in this section we use the suffix BL for objects mentioned in one of our blow-up lemmas which we will alter for the proof, following the style of the previous sentence.

For convenience, we shall collect the various conditions on the graphs H and G that we assume in our blow-up lemmas and those extra conditions that we derive in this section in a *general setup* (see Definition 2.23) which we shall use in many forthcoming lemmas. Moreover, we will define the concept of *good partial embedding* (see Definition 2.24), which is a collection of invariants we maintain in our proofs when embedding H.

This section is structured as follows. In Section 2.3.1 we collect and describe, mainly as reference for the reader, the various constants that appear in our blow-up lemma proofs. Sections 2.3.2 and 2.3.3 then list the properties we require of the refined partitions of H and G, respectively. In Section 2.3.4 we prove that partitions with these desired properties can be obtained from partitions supplied to our blow-up lemmas. In Section 2.3.5 we define some further key concepts used in our embedding procedures, such as candidate sets, available candidate sets and good partial embeddings, which then allows us to formally define our general setup. Finally, in Section 2.3.6, we define the notion of bad vertices which will, as explained in the proof overview, be vertices avoided in our embedding process, and prove Lemma 2.27 which states that most vertices in candidate sets are not bad.

2.3.1 Constants

In the following chapters, sections and lemmas we will use various constants. These are listed below, with their meaning. Firstly, we have the following constants which are chosen by the user of our blow-up lemmas.

Δ is the maximum degree of H.

Δ_J is the maximum number of times a vertex of Γ is a restricting vertex, that is, appears in a set J_x for $x \in V(H)$.

D is a constant governed by the degeneracy of H, which is only used in the degenerate graphs blow-up lemma.

d is the minimum p-density of regular pairs.

The user of our blow-up lemmas further chooses the constants $\Delta_{R'}^{\mathrm{BL}}$, α^{BL}, ζ^{BL}, κ^{BL} and r_1^{BL} (the last after being given $\varepsilon^{\mathrm{BL}}$ and ρ^{BL}). However, our preprocessing in Lemma 2.22 changes the values of these constants (and the reduced graphs R and R'), giving the following constants which we use in the rest of the paper.

$\Delta_{R'}$ is the maximum degree of the reduced graph R', which captures super-regular pairs.

α is the fraction of a part required to be potential buffer vertices.

ζ is the minimum relative size of any image restriction set I_x.

κ is the balancing factor (greater than 1), bounding $|V_i|/|V_j|$ for any pair of clusters of G (and hence parts of H).

r_1 is the upper bound on the number of clusters.

The blow-up lemmas guarantee the existence of the following constants.

$\varepsilon^{\mathrm{BL}}$ is the regularity we require in the regular partition provided by the user of the blow-up lemma.

ρ^{BL} is the fraction of vertices in each part of the partition of H supplied by the user which may be image restricted.

C is large, only appears in the random graphs blow-up lemmas and is the constant factor in the bound on the probability p.

c is small, only appears in the bijumbled graphs blow-up lemma and is the constant factor in the bound on the bijumbledness error term β.

Furthermore, the most important additional constants appearing in the proofs of the blow-up lemmas are the following.

μ is the fraction of each cluster of G contained in each of the "small" sets of the partition of G, thus $|V_i^{\mathrm{q}}| = |V_i^{\mathrm{c}}| = |V_i^{\mathrm{buf}}| = \mu|V_i|$ for each i.

ϑ comes into the maximum fraction ρp^ϑ of each part of H which may be image restricted. It is equal to either zero (proving Lemmas 1.21 and 1.23) or Δ (proving Lemma 1.25).

ρ appears in the maximum fraction ρp^ϑ of each part of H which may be image restricted. We also use ρ for a second purpose (which does not conflict with the first one): it is the fraction of non-image restricted vertices in X_i which may enter the queue during the RGA without causing the RGA to fail. Hence in total at most $2\rho|X_i|$ vertices of X_i may enter the queue during the RGA.

ε' is the global worst case regularity appearing in the proof of the blow-up lemmas. That is, whenever we use the fact that a pair is $(\cdot, d, p)$-regular in G, it will be at worst (ε', d, p)-regular.

$\varepsilon_{a,b}$ is the worst case regularity in the proofs between underlying restriction sets of adjacent vertices in H with respectively a and b previously embedded neighbours (see Section 2.3.5 for definitions of these terms).

ε is the initial regularity after pre-processing, and also controls the deterministic properties we require of Γ.

In order to make our proofs work it is enough to have constants in the size order

$$0 < \varepsilon \ll \varepsilon_{a,b} \le \varepsilon' \ll \rho \ll \mu \ll \alpha, \zeta, d, \Delta^{-1}, \Delta_{R'}^{-1}, \Delta_J^{-1}, D^{-1}, \kappa^{-1}$$
$$\text{and} \quad 0 < c, C^{-1} \ll r_1^{-1}, \varepsilon$$

where by $x \ll y$ we mean that there is a non-decreasing function $f \colon (0,1] \to (0,1]$ such that our proof works if $0 < x \le f(y)$. Observe that the constants on the right hand sides are (effectively) chosen by

the user of the blow-up lemma. We remark that one can then safely read this paper assuming that, for example, any function of ρ appearing in the proofs tending to zero with ρ is much smaller than any function of μ appearing in the proofs tending to zero with μ. However, for the convenience of the reader wishing to verify the proofs, we specify our constants (more or less) explicitly in each of the following results. We have made no attempt to optimise the values we give.

At the beginning of the proof of each of our blow-up lemmas we give explicit values for these various constants. In the proofs of the auxiliary lemmas, we begin with 'we require' various inequalities to hold. By this we mean that if these inequalities hold then the proof of the auxiliary lemma goes through, and it is easy to check that the required inequalities hold for our explicit values.

2.3.2 The partition of H

We shall refine the partition $\mathfrak{X}^{\mathrm{BL}}$ of H given as input to one of our blow-up lemmas to obtain a partition $\mathfrak{X} = \{X_i\}_{i\in[r]}$ of $V(H)$ with some additional properties, which we shall need in our proofs. We shall also refine the given partition of G (see Section 2.3.3 for the properties we require from this refined partition of G). In this subsection we will define the properties that we require from the refined partition of H. To put this definition into context (and provide some explanations on how it fits to the refinement of the partition of G in the subsequent subsection) we first need some explanations concerning our strategy in the proofs of the blow-up lemmas.

We shall (in a series of steps, some of which are performed in this preprocessing step, and some of which are performed on later parts of the proof) construct subsets X_i^{main}, X_i^{q}, X_i^{c} and X_i^{buf} of X_i. We define

$$X^{\mathrm{main}} = \bigcup_{i\in[r]} X_i^{\mathrm{main}}, \quad X^{\mathrm{q}} = \bigcup_{i\in[r]} X_i^{\mathrm{q}}, \quad X^{\mathrm{c}} = \bigcup_{i\in[r]} X_i^{\mathrm{c}}, \quad X^{\mathrm{buf}} = \bigcup_{i\in[r]} X_i^{\mathrm{buf}}.$$

Here, $X_i^{\mathrm{buf}} \subseteq \tilde{X}_i$ contains the *buffer vertices*, which will be chosen from the set of potential buffer vertices (see Section 2.3.4). The set X^{c} is only used in the proof of the random graphs blow-up lemma, Lemma 1.21; in the other proofs we set $X^{\mathrm{c}} = \emptyset$. The vertices in X^{c} shall form the *reserved cliques* which we will use to fix so-called *buffer defects* (see Section 3.4 for more details). We choose this set of vertices at the beginning of the proof of Lemma 1.21. Once we chose X^{buf} and X^{c}, we set $X_i^{\mathrm{main}} := X_i \setminus (X_i^{\mathrm{buf}} \cup X^{\mathrm{c}})$. These sets will have sizes $|X_i^{\mathrm{buf}}| = 4\mu|X_i|$ and $|X_i^{\mathrm{c}}| \leq \mu|X_i|$ and hence $(1-5\mu)|X_i| \leq |X_i^{\mathrm{main}}| \leq (1-4\mu)|X_i|$. The set $X^{\mathrm{q}} \subseteq X^{\mathrm{main}}$ will form the *queue* and only gets chosen during the embedding, in the random greedy algorithm; it will also be of size at most $\mu|X_i|$. We remark that in the following definition we shall only refer to the sets X_i and X_i^{buf}. We only chose to mention the sets X_i^{q} and X_i^{c} here as well because we need them to motivate the refinement of the partition of G in the next subsection.

The properties that we require of our refined partition and the buffer vertices are collected in the following definition. We remark that the reduced graphs R and R' used for this refined partition are blow-ups of the reduced graphs R_{BL} and R'_{BL} provided to the respective blow-up lemma. These properties bound the number of buffer vertices per part, and state that the first and second neighbourhoods of buffer vertices lie along edges of R'. They also assert that any pair of vertices within one part X_i, but also any pair of buffer vertices, and any pair consisting of a buffer vertex and an image restricted vertex, are far apart. Note that this implies that $H[X_i, X_j]$ forms a matching for any pair of parts X_i and X_j, and it also implies that buffer vertices may not be image restricted. We further require that all vertices in any

one X_i^{buf} have the same degree, and distinguish between clique buffers and non-clique buffers. The latter will be important for the selection of reserved cliques in the proof of Lemma 1.21 (Section 3.4).

Definition 2.20 (Good H-partition). We say that a partition $\mathcal{X} = \{X_i\}_{i \in [r]}$ of $V(H)$ is a *good H-partition* for reduced graphs $R' \subseteq R$ on vertex sets $[r]$, with buffer $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$, buffer vertices $X^{\text{buf}} \subseteq \bigcup_{i \in [r]} \tilde{X}_i$ and image restricting vertices $\mathcal{J} = \{J_x\}_{x \in V(H)}$, if the following conditions are satisfied for each $i, j, k \in [r]$.

- (H1) $(H, \mathcal{X})$ is an R-partition and $\tilde{\mathcal{X}}$ is an (α, R')-buffer for $(H, \mathcal{X})$.
- (H2) $\text{dist}_H(x, y) \geq 10$ and $J_x \cap J_y = \emptyset$ for each $x, y \in X_i$ with $x \neq y$.
- (H3) $|X_i^{\text{buf}}| = 4\mu |X_i|$, and $|\{x \in X_i : x \in N(X^{\text{buf}})\}| \leq 4\kappa \Delta_{R'} \mu |X_i|$.
- (H4) $\text{dist}_H(x, y) \geq 5$ for each $x, y \in X^{\text{buf}}$ with $x \neq y$.
- (H5) $\text{dist}_H(x, y) \geq 3$ for each $x \in X^{\text{buf}}$ and each image restricted y.
- (H6) All vertices of X_i^{buf} have degree b for some $0 \leq b \leq \Delta$. We then call X_i^{buf} a *degree-b buffer*.
- (H7) Either all or none of the vertices of X_i^{buf} are in copies of $K_{\Delta+1}$. We then call X_i^{buf} a *clique buffer* or *non-clique buffer*, respectively.
- (H8) If X_i^{buf} is a clique buffer, then at least $\frac{1}{2\Delta+4}\alpha |X_i|$ vertices of $\tilde{X}_i$ are in copies of $K_{\Delta+1}$ which do not contain either vertices of X^{buf} or image restricted vertices.

Note that in this definition we do need to know about image restricting vertices (since they play a rôle in (H2)), which are in Γ, and we do need to know which vertices of H are image restricted (since they play a rôle in (H5)), but we do not need to know anything about G.

2.3.3 The partition of G

We next define the properties we require from our refined partition of G. Firstly, we shall subdivide each cluster of the partition $\mathcal{V}^{\text{BL}}$ of G, which is given as input to one of our blow-up lemmas, into several new clusters to obtain a new partition $\mathcal{V}$ matching the refined partition of H. Further, as also mentioned in the proof overview, we shall partition each new cluster V_i into parts V_i^{main}, V_i^{q}, V_i^{c} and V_i^{buf}. Again, we set

$$V^{\text{main}} = \bigcup_{i \in [r]} V_i^{\text{main}}, \quad V^{\text{q}} = \bigcup_{i \in [r]} V_i^{\text{q}}, \quad V^{\text{c}} = \bigcup_{i \in [r]} V_i^{\text{c}}, \quad V^{\text{buf}} = \bigcup_{i \in [r]} V_i^{\text{buf}}.$$

As previously indicated, each vertex $x \in V(H)$ will be embedded into the cluster V_i such that $x \in X_i$. In the rest of the paper, we often refer to V_i simply by $V(x)$, to V_i^{main} by $V^{\text{main}}(x)$, and so on, to avoid the use of indices whenever these are not important.

As explained before, most of the vertices from V^{main} will be used to embed most vertices of X^{main}, while some vertices from V^{q} will be used to embed the remainder, which constitutes the queue X^{q}. We also have the set V^{c}, the *clique reservoir*, that we will use in the proof of the random graphs blow-up lemma to embed some of the vertices in X^{c}. Some vertices of each X_i^{c} may need to be used to fix so-called buffer defects, and these could be embedded anywhere in V_i. All other vertices of each X_i^{c} get embedded to V_i^{c}. Finally, the vertices from V^{buf} together with all remaining vertices of $V^{\text{main}} \cup V^{\text{q}} \cup V^{\text{c}}$, that is, vertices that were not used when embedding $X^{\text{main}} \cup X^{\text{q}} \cup X^{\text{c}}$, will be used for embedding the buffer vertices X^{buf}.

The properties we require of a good G-partition are as follows. It must be a regular R-partition, and furthermore the one-sided and two-sided inheritance properties of $\mathcal{V}^{\mathrm{BL}}$ on R'_{BL} must be transferred to $\mathcal{V}$ on R'. The super-regularity properties of $\mathcal{V}^{\mathrm{BL}}$ on R'_{BL} must be transferred to give minimum degree conditions for each of the parts V_i^{main}, V_i^{q} and V_i^{c}. Finally, we need to create new image restriction sets $\mathcal{I}$ from the supplied $\mathcal{I}^{\mathrm{BL}}$ which with the (unchanged) $\mathcal{J}$ form a restriction pair.

Definition 2.21 (Good G-partition). Given a good H-partition of $V(H)$ for reduced graphs $R' \subseteq R$ on vertex sets $[r]$, with buffer $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ and buffer vertices $X^{\mathrm{buf}} \subseteq \bigcup_{i \in [r]} \tilde{X}_i$, with image restrictions $\mathcal{I} = \{I_x\}_{x \in V(H)}$, and restricting vertices $\mathcal{J} = \{J_x\}_{x \in V(H)}$, we say that a partition $\mathcal{V} = \{V_i\}_{i \in [r]}$ of $V(G)$ with a partition $V_i = V_i^{\mathrm{main}} \dot\cup V_i^{\mathrm{q}} \dot\cup V_i^{\mathrm{c}} \dot\cup V_i^{\mathrm{buf}}$ of each cluster is a *good G-partition* for $R' \subseteq R$, if the following conditions are satisfied.

(G1) For each i we have $|V_i^{\mathrm{main}}| = (1 - 3\mu)|V_i|$ and $|V_i^{\mathrm{q}}| = |V_i^{\mathrm{c}}| = |V_i^{\mathrm{buf}}| = \mu|V_i|$.

(G2) $(G, \mathcal{V})$ is an (ε, d, p)-regular R-partition, which has one-sided inheritance on R', and two-sided inheritance on R' for $\tilde{\mathcal{X}}$.

(G3) For each $ij \in R'$ and $v \in V_i$, we have

$$\deg_G(v; V_j^{\mathrm{main}}) \geq (1 - 3\mu)(d - \varepsilon) \max\left(p|V_j|, \deg_\Gamma(v; V_j)/2\right),$$
$$\deg_G(v; V_j^{\mathrm{q}}) \geq \mu(d - \varepsilon) \max\left(p|V_j|, \deg_\Gamma(v; V_j)/2\right),$$
$$\deg_G(v; V_j^{\mathrm{c}}) \geq \mu(d - \varepsilon) \max\left(p|V_j|, \deg_\Gamma(v; V_j)/2\right).$$

(G4) For each $x \in X_i$, we have

$$|I_x \cap V_i^{\mathrm{main}}| \geq (1 - \varepsilon)(1 - 3\mu)|I_x|,$$
$$|I_x \cap V_i^{\mathrm{q}}|, |I_x \cap V_i^{\mathrm{c}}|, |I_x \cap V_i^{\mathrm{buf}}| \geq (1 - \varepsilon)\mu|I_x|,$$
$$\left|N^*(J_x; V_i^{\mathrm{main}})\right| = (1 \pm \varepsilon)(1 - 3\mu)p^{|J_x|}|V_i|, \text{ and}$$
$$\left|N^*(J_x; V_i^{\mathrm{q}})\right|, \left|N^*(J_x; V_i^{\mathrm{c}})\right|, \left|N^*(J_x; V_i^{\mathrm{buf}})\right| = (1 \pm \varepsilon)\mu p^{|J_x|}|V_i|.$$

(G5) $\mathcal{I}$ and $\mathcal{J}$ form a $(\rho p^\vartheta, \zeta, \Delta, \Delta_J)$-restriction pair for the partitions $\mathcal{X}$ and $\mathcal{V}$.

(G6) For each i, the number of image restricted vertices in X_i is at most $\rho p^\vartheta |X_i|$.

Note that the sizes of the sets V_i^{main}, V_i^{q}, V_i^{c}, V_i^{buf} are *not* the same as the sizes of the sets X_i^{main}, X_i^{q}, X_i^{c} and X_i^{buf}. For example we will have $|X_i^{\mathrm{main}}| = (1 - 4\mu)|X_i| < (1 - 3\mu)|V_i| = |V_i^{\mathrm{main}}|$, and $|X_i^{\mathrm{q}}|, |X_i^{\mathrm{c}}| \leq 2\rho|X_i|$, which is much smaller than $\mu|V_i|$. Thus, the subsets X_i^{main}, X_i^{q}, X_i^{c} of X_i will be significantly smaller than the corresponding sets in V_i, while the set X_i^{buf} is larger than V_i^{buf}. We also note that we do require two-sided inheritance for $\tilde{\mathcal{X}}$ in (G2), not just for X^{buf}. In the proof of Lemma 1.21, when dealing with buffer defects, we will make use of this stronger statement.

2.3.4 Obtaining a good H-partition and a good G-partition

The following lemma proves that we can obtain good partitions of H and G from the partitions provided to our blow-up lemmas. The proof of this lemma is straightforward, though not short since several conditions must be checked. Briefly, the idea is that we will draw an auxiliary graph F_i on each part

X_i^{BL} of $\mathfrak{X}^{\mathrm{BL}}$ with edges joining pairs of vertices which are at distance less than 10 or share an image restricting vertex. We will apply our modification of the Hajnal-Szemerédi theorem (Lemma 2.4) to each of these graphs together with the set $\tilde{X}_i$ to obtain a partition $\mathfrak{X}$ and reduced graphs R and R' satisfying the first three conditions of a good H-partition, and then construct the sets X_i^{buf} greedily to obtain the remaining properties. We will then randomly refine the partition $\mathcal{V}^{\mathrm{BL}}$ to obtain a matching partition $\mathcal{V}$, and further randomly split each part V_i of $\mathcal{V}$ into V_i^{main}, V_i^{q}, V_i^{c} and V_i^{buf}. We will use the concentration inequalities in Theorem 2.1 to obtain concentration for various set sizes in the randomly chosen parts, which in particular show that with high probability we obtain the desired good G-partition. We stress that although p with a lower bound reminiscent of our random graph blow-up lemmas makes an appearance in the lemma statement, this does not mean we are going to assume Γ is a random graph. We will simply need to know that the quantity $p^b n$ is large compared to $\log n$.

Lemma 2.22 (Good partitions lemma). *For all positive integers* Δ, $\Delta_{R'}$, Δ_J, r_1, ϑ *and* b, *all* $\kappa > 2$, *all* $\alpha, \zeta, d, \varepsilon > 0$, *and all sufficiently small* $\mu, \rho > 0$, *there exists* C *such that the following holds whenever* $p \geq C\left(\frac{\log n}{n}\right)^{1/b}$ *and* $n > C$. *Let* $\delta = \frac{1}{8}(\Delta + \Delta_J)^{-10}$ *and* $r \leq r_1$.
Let R_{BL} *be a graph on* $r_{\mathrm{BL}} = \delta r$ *vertices and let* $R'_{\mathrm{BL}} \subseteq R_{\mathrm{BL}}$ *be a spanning subgraph with* $\Delta(R'_{\mathrm{BL}}) \leq \Delta_{R'}^{\mathrm{BL}} = \delta \Delta_{R'}$. *Let* H *and* $G \subseteq \Gamma$ *be graphs with* $\frac{1}{2}\kappa$-*balanced size-compatible vertex partitions* $\mathfrak{X}^{\mathrm{BL}} = \{X_i^{\mathrm{BL}}\}_{i \in [r_{\mathrm{BL}}]}$ *and* $\mathcal{V}^{\mathrm{BL}} = \{V_i^{\mathrm{BL}}\}_{i \in [r_{\mathrm{BL}}]}$, *respectively, which have parts of size at least* $2n/(\kappa \delta r_1)$. *Let* $\tilde{\mathfrak{X}}^{\mathrm{BL}} = \{\tilde{X}_i^{\mathrm{BL}}\}_{i \in [r_{\mathrm{BL}}]}$ *be a family of subsets of* $V(H)$, $\mathfrak{I}^{\mathrm{BL}} = \{I_x^{\mathrm{BL}}\}_{x \in V(H)}$ *be a family of image restrictions, and* $\mathfrak{J} = \{J_x\}_{x \in V(H)}$ *be a family of restricting vertices. Suppose that*

(i) $\Delta(H) \leq \Delta$, $(H, \mathfrak{X}^{\mathrm{BL}})$ *is an* R_{BL}-*partition, and* $\tilde{\mathfrak{X}}^{\mathrm{BL}} = \{\tilde{X}_i^{\mathrm{BL}}\}_{i \in [r_{\mathrm{BL}}]}$ *is a* $(2\alpha, R'_{\mathrm{BL}})$-*buffer for* $(H, \mathfrak{X})$,

(ii) $(G, \mathcal{V}^{\mathrm{BL}})$ *is a* $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-*regular* R_{BL}-*partition, which is* $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-*super-regular on* R'_{BL}, *has one-sided inheritance on* R', *and two-sided inheritance on* R' *for* $\tilde{\mathfrak{X}}^{\mathrm{BL}}$,

(iii) $\mathfrak{I}^{\mathrm{BL}}$ *and* $\mathfrak{J}$ *form a* $\left(\frac{1}{2}\delta\rho p^\vartheta, 2\zeta, \Delta, \Delta_J\right)$-*restriction pair*, *and* $|J_x| \leq b$ *for each* $x \in V(H)$.

Then there is a graph R *on* r *vertices and a spanning subgraph* $R' \subseteq R$ *with* $\Delta(R') \leq \Delta_{R'}$, *together with* κ-*balanced size-compatible partitions* $\mathfrak{X} = \{X_i\}_{i \in [r]}$ *of* H *and* $\mathcal{V} = \{V_i\}_{i \in [r]}$ *of* G, *which have parts of size at least* $n/(\kappa r_1)$, *a family* $\tilde{\mathfrak{X}} = \{\tilde{X}_i\}_{i \in [r]}$ *of potential buffer vertices, a family* $\mathfrak{I} = \{I_x\}_{x \in V(H)}$ *of image restrictions, subsets* $X_i^{\mathrm{buf}} \subseteq \tilde{X}_i$ *for each* $i \in [r]$, *and a partition* $V_i = V_i^{\mathrm{main}} \, \dot\cup \, V_i^{\mathrm{q}} \, \dot\cup \, V_i^{\mathrm{c}} \, \dot\cup \, V_i^{\mathrm{buf}}$ *for each* $i \in [r]$, *which give a good* H-*partition and a corresponding good* G-*partition.*

Proof. Let $\delta = \frac{1}{8}(\Delta + \Delta_J)^{-10}$. We set $C = 10^6 \delta^{-1} \mu^{-1} \varepsilon^{-2} d^{-\Delta} \zeta^{-1} \kappa r_1$. We require

$$\mu < \frac{\alpha}{10000 \kappa \Delta_{R'}^4 \Delta^{10} (\Delta + 2)} \qquad \text{and} \qquad \rho \leq \mu.$$

First we refine the vertex partition $\mathfrak{X}^{\mathrm{BL}}$ of H to obtain properties $(H1)$ and $(H2)$. For each $i \in [r_{\mathrm{BL}}]$ we define a graph F_i on the vertex set X_i^{BL} by putting an edge between x and x' whenever either $J_x \cap J_{x'} \neq \emptyset$ or x and x' are at distance less than 10 in H. Observe that $\Delta(F_i) \leq (\Delta + \Delta_J)^{10}$. We now apply Lemma 2.4 to F_i with the set $\tilde{X}_i^{\mathrm{BL}}$. This returns an equitable partition of the vertices X_i^{BL} into independent sets in F_i which also partitions $\tilde{X}_i^{\mathrm{BL}}$ equitably. We do this for each $i \in [r_{\mathrm{BL}}]$ to obtain a partition $\mathfrak{X}$ of $V(H)$ into $8(\Delta + \Delta_J)^{10} r_{\mathrm{BL}} = \delta^{-1} r_{\mathrm{BL}} = r$ parts and a family $\tilde{\mathfrak{X}} = \{\tilde{X}_i\}_{i \in [r]}$. We let R be the graph obtained from R_{BL} by replacing each vertex with an independent set on $8(\Delta + \Delta_J)^{10} = \delta^{-1}$ vertices, and each edge with a

complete bipartite graph between the corresponding sets. We obtain R' similarly from R'_{BL}. Thus $(H, \mathcal{X})$ is an R-partition, satisfying $(H2)$ by construction. Furthermore, we have $\Delta(R') = \delta^{-1}\Delta(R'_{\mathrm{BL}}) \le \Delta_{R'}$.

Since each X_i^{BL} with $i \in [r_{\mathrm{BL}}]$ has size at least $2n/(\kappa \delta r_1)$ and is equipartitioned into δ^{-1} parts, we see that, because n is chosen sufficiently large, we have $|X_i| \ge n/(\kappa r_1)$ for each $i \in [r]$. Given $i, j \in [r]$ let $i', j' \in [r_{\mathrm{BL}}]$ be such that $X_i \subseteq X_{i'}^{\mathrm{BL}}$ and $X_j \subseteq X_{j'}^{\mathrm{BL}}$. Then we have

$$|X_i| \le \delta|X_{i'}^{\mathrm{BL}}| + 1 \le \tfrac{\kappa}{2}\delta|X_{j'}^{\mathrm{BL}}| + 1 \le \tfrac{\kappa}{2}(|X_j| + 1) + 1 \le \kappa|X_j|$$

where the final inequality is since $|X_j| \ge n/(\kappa r_1)$ and n is sufficiently large. Thus $\mathcal{X}$ is κ-balanced.

Given $i \in [r]$ let $i' \in [r_{\mathrm{BL}}]$ be such that $X_i \subseteq X_{i'}^{\mathrm{BL}}$. Then we have

$$|\tilde{X}_i| \ge \delta|\tilde{X}_{i'}^{\mathrm{BL}}| - 1 \ge 2\alpha\delta|X_{i'}^{\mathrm{BL}}| - 1 \ge 2\alpha(|X_i| - 1) - 1 \ge \alpha|X_i|$$

where again the final inequality is since n is sufficiently large. Thus $\tilde{\mathcal{X}}$ forms an (α, R')-buffer for $(H, \mathcal{X})$.

Before we verify the remaining good H-partition properties, it is convenient to check $(G6)$. By assumption (iii) of the lemma, at most $\frac{1}{2}\delta\rho p^{\vartheta}|X_{i'}^{\mathrm{BL}}|$ vertices of any given $X_{i'}^{\mathrm{BL}}$ are image restricted. Thus for any i such that $X_i \subseteq X_{i'}^{\mathrm{BL}}$, at most

$$\tfrac{1}{2}\delta\rho p^{\vartheta}|X_{i'}^{\mathrm{BL}}| \le \tfrac{1}{2}\rho p^{\vartheta}\left(|X_i| + 1\right) \le \rho p^{\vartheta}|X_i|$$

vertices of X_i are image restricted, as required for $(G6)$.

We proceed with properties $(H3)$–$(H8)$ by choosing buffer vertices X_i^{buf} from the potential buffer vertices $\tilde{X}_i$ for each $i \in [r]$ sequentially. For a given $i \in [r]$, we first decide on the degree of the buffer and whether or not it is a clique buffer. We split the vertices of $\tilde{X}_i$ into $\Delta + 2$ subsets, one for each of the possible degrees in $\{0, \ldots, \Delta - 1\}$, one for vertices of degree Δ not in copies of $K_{\Delta+1}$, and one for vertices in copies of $K_{\Delta+1}$. We take a largest one S_i of these sets, and will choose X_i^{buf} within it. We do this greedily, one vertex at a time, always picking a vertex at distance at least five from previously chosen buffer vertices, and at distance at least three from any image restricted vertices. We now justify that it is possible to do this, that is, that we can pick the desired $4\mu|X_i|$ vertices in S_i without running out of vertices in S_i.

Since all neighbours (respectively, second neighbours) in H of vertices in $\tilde{X}_i$ are in clusters V_j with $ij \in R'$ (respectively, V_k with $ij, jk \in R'$), by $(G6)$, verified above, it follows that the number of image restricted vertices at distance 2 or less to a vertex of $\tilde{X}_i$ is at most $(\Delta_{R'}^2 + \Delta_{R'} + 1)\rho\kappa|X_i|$. Each of these vertices has at most one neighbour or second neighbour in $\tilde{X}_i'$, since the distance between any two vertices of X_i is at least ten. Similarly, if x is in some $\tilde{X}_k$ and is at distance four or less from some vertex of $\tilde{X}_i$, then the path between x and $\tilde{X}_i$ follows edges of R', and there is only one vertex of $\tilde{X}_i$ which is at distance four or less from x. Thus the total number of vertices of $\tilde{X}_i$ which are at distance four or less from some vertex of X^{buf} is at most

$$\left(1 + \Delta_{R'} + \Delta_{R'}^2 + \Delta_{R'}^3 + \Delta_{R'}^4\right) \cdot 4\mu\kappa|X_i|.$$

Since $|S_i| \ge \frac{1}{\Delta+2}|\tilde{X}_i|$, it follows that at each step the number of vertices in $S_i \subseteq \tilde{X}_i$ we have to choose from is at least

$$\tfrac{1}{\Delta+2}\alpha|X_i| - 3\Delta_{R'}^2\rho\kappa|X_i| - 20\Delta_{R'}^4\mu\kappa|X_i|,$$

which is greater than $\frac{1}{2\Delta+4}\alpha|X_i|$ by choice of μ and ρ. Thus we can always choose the desired vertices of X_i^{buf}, giving $(H4)$–$(H7)$. In the event that X_i^{buf} is a clique buffer, we have at the end at least $\frac{1}{2\Delta+4}\alpha|X_i|$

vertices of $\tilde{X}_i$ remaining which are in copies of $K_{\Delta+1}$ (by definition of S_i) that do not contain either vertices in X^{buf} or image restricted vertices, giving $(H\,8)$.

Finally, neighbours of X^{buf} in X_i must be in parts X_j such that $ij \in R'$. Since any two vertices in X_i are at distance at least ten, each such vertex has at most one neighbour in X_i and so at most $\Delta_{R'} 4\mu\kappa |X_i|$ vertices of X_i are in $N(X^{\mathrm{buf}})$, completing $(H\,3)$. This establishes that we have a good H-partition.

We now construct a corresponding good G-partition. The construction is simple: for each $V_i^{\mathrm{BL}} \in \mathcal{V}^{\mathrm{BL}}$ we choose, uniformly at random, an equitable partition into δ^{-1} parts, which parts we associate (arbitrarily) to subparts of X_i^{BL} of corresponding size, obtaining $\mathcal{V}$. Note that this is possible since $|X_i^{\mathrm{BL}}| = |V_i^{\mathrm{BL}}|$. We then, for each $i \in [r]$, choose uniformly at random a partition of V_i into one set of size $(1-3\mu)|V_i|$ and three of size $\mu|V_i|$, which form V_i^{main}, V_i^{q}, V_i^{c} and V_i^{buf} respectively. For each $x \in V_i$, we let $I_x = I_x^{\mathrm{BL}} \cap V_i$.

By construction $\mathcal{X}$ and $\mathcal{V}$ are size-compatible, so what remains is to show that $\mathcal{V}$ is with positive probability a good G-partition. In fact, we will show this holds asymptotically almost surely. Recall that we have already established $(G\,6)$, and by construction $(G\,1)$ holds, so that the remaining conditions are $(G\,2)$–$(G\,5)$.

Observe that the sets V_i, V_i^{main}, V_i^{q}, V_i^{c} and V_i^{buf} are all distributed as the uniform random set of that size within $V_{i'}^{\mathrm{BL}}$, where $V_i \subseteq V_{i'}^{\mathrm{BL}}$, though they are of course not independent. It follows that for any set $U \subseteq V_{i'}^{\mathrm{BL}}$, the sizes of each of the sets $V_i \cap U$, $V_i^{\mathrm{main}} \cap U$, $V_i^{\mathrm{q}} \cap U$, $V_i^{\mathrm{c}} \cap U$ and $V_i^{\mathrm{buf}} \cap U$ are hypergeometrically distributed, each with (since n is sufficiently large) expectation at least $\frac{1}{2}\mu\delta|U|$. In particular, if $|U| \geq 10^3\varepsilon^{-2}\log n$, then the probability that any given one of these sizes is not within a $\left(1 \pm \frac{\varepsilon}{10}\right)$-factor of its expectation is by Theorem 2.1 at most

$$2e^{-\varepsilon^2|U|/300} \leq n^{-2}\,.$$

We now give various sets U to which we will apply this observation, and verify that each is sufficiently large. For each $ij \in R'$ and each $v \in V_j$, by construction of R' and since $(G, \mathcal{V}^{\mathrm{BL}})$ is super-regular on R'_{BL}, we have

$$\begin{aligned}
\deg_\Gamma(v; V_{i'}^{\mathrm{BL}}) &\geq \deg_G(v; V_{i'}^{\mathrm{BL}}) \\
&\geq (d-\varepsilon)\max\left\{p|V_{i'}^{\mathrm{BL}}|, \tfrac{1}{2}\deg_\Gamma(v; V_{i'}^{\mathrm{BL}})\right\} \geq 10^3\varepsilon^{-2}\log n\,,
\end{aligned}$$

where the final inequality holds by choice of p and C. We can thus take U to be $N_\Gamma(v; V_{i'}^{\mathrm{BL}})$ or $N_G(v; V_{i'}^{\mathrm{BL}})$.

For each $x \in X_i$, where $X_i \subseteq X_{i'}^{\mathrm{BL}}$, since $\left(\mathcal{J}^{\mathrm{BL}}, \mathcal{J}\right)$ form a $\left(\frac{1}{2}\delta\rho, 2\zeta, \Delta, \Delta_J\right)$-restriction pair, we have

$$\left|N_\Gamma^*(J_x; V_{i'}^{\mathrm{BL}})\right| \geq |I_x^{\mathrm{BL}}| \geq \zeta(dp)^{|J_x|}|V_{i'}^{\mathrm{BL}}| \geq 10^3\varepsilon^{-2}\log n\,,$$

where the final inequality comes from the assumption $|J_x| \leq b$ and our choice of p and C. We can thus also take U to be $N_\Gamma^*(J_x; V_{i'}^{\mathrm{BL}})$ or I_x.

In total, we have at most $5r$ randomly chosen sets in which we are interested, and at most $4n$ choices of U with which we intersect the random sets. Taking the union bound, we see that a.a.s. each of these intersections has size within a $\left(1 \pm \frac{\varepsilon}{10}\right)$-factor of its expectation. In particular, there exists a partition $\mathcal{V}$ in which each of these intersections is within a $\left(1 \pm \frac{\varepsilon}{10}\right)$-factor of its expectation. We fix such a partition, and prove that it is the desired good G-partition.

We begin with $(G\,2)$. Suppose $ij \in R$. Then V_i and V_j come from two parts $V_{i'}^{\mathrm{BL}}$ and $V_{j'}^{\mathrm{BL}}$ of $\mathcal{V}^{\mathrm{BL}}$ which form a $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-regular pair, and we have $|V_i| > \frac{1}{2}\delta|V_{i'}^{\mathrm{BL}}|$ and $|V_j| > \frac{1}{2}\delta|V_{j'}^{\mathrm{BL}}|$. It follows from the definition that (V_i, V_j) is (ε, d, p)-regular, so $(G, \mathcal{V})$ forms an (ε, d, p)-regular R-partition. Given now $ij, jk \in R'$, with $V_j \subseteq V_j^{\mathrm{BL}}$ and $V_k \subseteq V_k^{\mathrm{BL}}$, and any $v \in V_i$, we have by construction of $\mathcal{V}$ that $\left||N_\Gamma(v; V_j)|\right| \geq \frac{1}{2}\delta\left|N_\Gamma(v; V_{j'})\right|$ and that $|V_k| \geq \frac{1}{2}\delta|V_{k'}|$. Since the pair $\left(N_\Gamma(v; V_{j'}), V_{k'}\right)$ is $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-regular in G, it follows from the definition that the pair $\left(N_\Gamma(v; V_j), V_k\right)$ is (ε, d, p)-regular in G, so that we have one-sided inheritance on R'. A similar argument shows that if also $jk \in R'$ and there is a triangle of H with one vertex in each of $\tilde{X}_i$, X_j and X_k, then also $\left(N_\Gamma(v; V_j), N_\Gamma(v; V_k)\right)$ is (ε, d, p)-regular in G, so that we have two-sided inheritance on R' for $\tilde{\mathcal{X}}$.

Both of $(G\,3)$ and $(G\,4)$ follow directly from the construction of $\mathcal{V}$ together with, respectively, the $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-super-regularity of $(G, \mathcal{V}^{\mathrm{BL}})$ on R'_{BL} and that $(\mathcal{J}^{\mathrm{BL}}, \mathcal{J})$ forms a $\left(\frac{1}{2}\delta\rho p\,p^{\vartheta}, 2\zeta, \Delta, \Delta_J\right)$-restriction pair. The same holds for everything except (f) in the definition of a restriction pair for $(G\,5)$, since $(G\,6)$ was already established. Given an image restricted $x \in X_i \subseteq X_{i'}^{\mathrm{BL}}$ and an edge $xy \in E(H)$ with $Y \in X_j \subseteq X_j^{\mathrm{BL}}$, the pair $\left(N_\Gamma^*(J_x; V_{i'}), N_\Gamma^*(J_y; V_{j'})\right)$ is $\left(\frac{1}{2}\delta\varepsilon, d, p\right)$-regular by assumption. Since $\left|N_\Gamma^*(J_x; V_i)\right| \geq \frac{1}{2}\delta\left|N_\Gamma^*(J_x; V_{i'})\right|$, and similarly for y, it follows from the definition that $\left(N_\Gamma^*(J_x; V_i), N_\Gamma^*(J_y; V_j)\right)$ is (ε, d, p)-regular, completing the verification of $(G\,5)$. $\qquad\square$

With this we are ready to define the general setup which is common to our proofs and encapsulates the definitions of good embeddings and the constant choices above. This encapsulation is convenient because we will use this general setup in many of the lemmas to come.

Definition 2.23 (General Setup). When we say that *we assume the General Setup* we assert that constant choices satisfying the conditions of Section 2.3.1 have been made, that there exist κ-balanced size-compatible partitions $\mathcal{X}$ and $\mathcal{V}$ of $V(H)$ and $V(G)$ respectively whose parts are of size at least $\frac{n}{\kappa r_1}$, together with reduced graphs R and R' with $\Delta(R') \leq \Delta_{R'}$, an image restriction pair $(\mathcal{I}, \mathcal{J})$, an (α, R')-buffer $\tilde{\mathcal{X}}$, buffer vertices X^{buf}, and sets V^{main}, V^{q}, V^{c} and V^{buf}, that the partition of $V(H)$ is a good H-partition, and that the partition of G is a good G-partition.

2.3.5 Underlying restrictions, candidates and good partial embeddings

In this section we provide definitions of a number of additional sets that will play important roles in our embedding algorithms. We also state which invariants are maintained throughout the embedding, and encapsulate this in the definition of good partial embeddings.

Suppose ψ is a *partial embedding* of H into G, that is, an injective graph homomorphism with domain $\mathrm{Dom}(\psi) \subseteq V(H)$. We call a vertex $x \in H$ *embedded* if $x \in \mathrm{Dom}(\psi)$, and otherwise *unembedded*. Of course this is with respect to ψ, but this will always be clear from the context and we will not in future mention it. The set of embedded neighbours of x is

$$\Pi(x) := \left\{\psi(y) : xy \in E(H), y \in \mathrm{Dom}(\psi)\right\}$$

and we define

$$\pi(x) := \left|\Pi(x)\right| \quad \text{and} \quad \pi^*(x) = \left|\Pi(x)\right| + \left|J_x\right|.$$
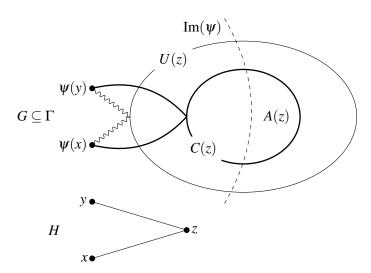
Figure 2.1: Underlying restriction sets, candidate sets, and available candidate sets for an example in which x and y are embedded and their common neighbour z is not. The thick lines are edges of G, and the wavy lines of Γ.

Thus, $\pi^*(x)$ counts the number of already embedded neighbours of x by the partial embedding ψ plus the 'neighbours' from J_x (if any).

Given an unembedded vertex $x \in X_i$, we define the *underlying restriction set* $U(x)$, the *candidate set* $C(x)$, and the *available candidate set* $A(x)$ of x by

$$U(x) := V(x) \cap N_\Gamma^*\big(\Pi(x) \cup J_x\big),$$
$$C(x) := I_x \cap N_G^*\big(\Pi(x)\big), \text{ and}$$
$$A(x) := C(x) \setminus \operatorname{Im}(\psi), \text{ respectively.}$$

For an illustration see Figure 2.1.

Recall that G is a subgraph of Γ (being either a random or a bijumbled graph), which is typically far from being complete. The underlying set signifies thus the vertices that could in principle be used for embedding (if $G = \Gamma$), whereas the candidate sets denote the vertices that are possible *in G*, and the available candidate sets only consider those vertices from the candidate sets which have not been previously taken as images by the vertices embedded earlier (in $\operatorname{Im}(\psi)$). We also refer to vertices from $C(x)$ as candidate vertices for x. Further, we define

$$U^{\text{main}}(x) := U(x) \cap V^{\text{main}}(x), \qquad\qquad U^{\text{q}}(x) := U(x) \cap V^{\text{q}}(x),$$
$$U^{\text{c}}(x) := U(x) \cap V^{\text{c}}(x), \qquad\qquad U^{\text{buf}}(x) := U(x) \cap V^{\text{buf}}(x),$$

and similarly for the respective subsets of the candidate set $C(x)$ and the available candidate set $A(x)$ of x. For an illustration see Figure 2.2.

Further we define the *underlying restriction graph* between $Y \subseteq X_i$ and $Z \subseteq V_i$ to be the bipartite graph with parts Y and Z and edges yz whenever $z \in U(y)$. We define similarly the *candidate graph*
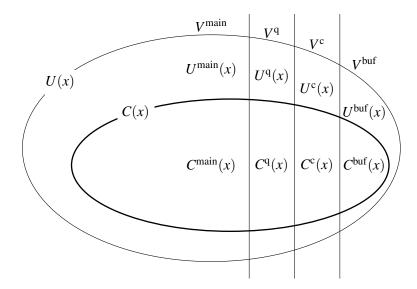
Figure 2.2: Partitioning underlying restriction sets and candidate sets into one big and three small parts

between $Y \subseteq X_i$ and $Z \subseteq V_i$ to be the bipartite graph with parts Y and Z and edges yz whenever $z \in C(y)$. If $x \in V(H)$ and $v \in C(x)$ we say v is a *candidate* for x.

Now we can introduce the concept of good partial embeddings. The intention of this definition is to fix properties which easily allow to prove that a good partial embedding is 'locally' good: If x is unembedded, then almost all of $C(x)$ consists of vertices v such that if v is not in $\text{Im}(\psi)$, then $\psi \cup \{x \to v\}$ is a good partial embedding (this assertion is made formal in Lemma 2.27). For an illustration see Figure 2.3.

Definition 2.24 (Good partial embedding). We call ψ a *good partial embedding* if the following conditions hold.

($GPE\,1$) For each $x \in \text{Dom}(\psi)$ we have $\psi(x) \in I_x$.

($GPE\,2$) For each unembedded $x \in X_i$ we have

$$\left|U(x)\right| = (p \pm \varepsilon p)^{\pi^*(x)}\left|V(x)\right|, \quad \left|U^{\text{main}}(x)\right| = (1-3\mu)(p \pm \varepsilon p)^{\pi^*(x)}\left|V(x)\right|, \quad \text{and}$$
$$\left|U^{\text{q}}(x)\right|, \left|U^{\text{c}}(x)\right|, \left|U^{\text{buf}}(x)\right| = \mu(p \pm \varepsilon p)^{\pi^*(x)}\left|V(x)\right|.$$

($GPE\,3$) For each unembedded x we have

$$\left|C(x)\right| \geq (1-\varepsilon')(dp - \varepsilon' p)^{\pi(x)}\left|I_x\right|,$$
$$\left|C^{\text{main}}(x)\right| \geq (1-\varepsilon')(1-3\mu)(dp - \varepsilon' p)^{\pi(x)}\left|I_x\right|, \quad \text{and}$$
$$\left|C^{\text{q}}(x)\right|, \left|C^{\text{c}}(x)\right|, \left|C^{\text{buf}}(x)\right| \geq (1-\varepsilon')\mu(dp - \varepsilon' p)^{\pi(x)}\left|I_x\right|.$$

($GPE\,4$) For each $xy \in E(H)$ with x and y unembedded, the pair $\left(U(x), U(y)\right)$ is $(\varepsilon_{\pi^*(x),\pi^*(y)}, d, p)$-regular in G.

It is important to observe that the trivial partial embedding, in which no vertices are embedded, is not automatically a good partial embedding: image restricted vertices might destroy any of ($GPE\,2$)–($GPE\,4$).
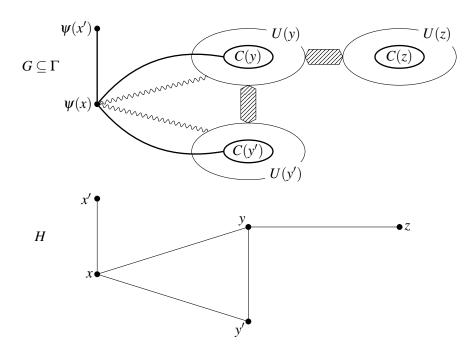
Figure 2.3: Good partial embeddings form embeddings when restricted to the mapped vertices x, x' and guarantee sizes and regularity for the underlying restriction sets of unmapped vertices y, y', z. The thick lines are edges of G, the wavy lines of Γ, and the hatched areas represent regular pairs in G.

However, it follows from straightforward checking of the definitions that if we are provided with a good H-partition and a corresponding good G-partition, then the trivial partial embedding is a good partial embedding.

Fact 2.25. *If we assume the General Setup for H and G then the trivial embedding is a good partial embedding of H into G.*

Finally, in much of the rest of this paper we will want to consider not just one partial embedding, but a sequence $\psi_0, \dots$ of them. We will want to refer to sets $A^{\mathrm{main}}(x)$, $C^{\mathrm{main}}(x)$ and so on, and quantities $\pi(x)$ and so on, with reference to each of these. We will do this by following the convention that a subscript t attached to any of these means it is with reference to ψ_t. Thus we write $A_t^{\mathrm{main}}(x)$ for the set $A^{\mathrm{main}}(x)$ with respect to ψ_t, and similarly $C_t^{\mathrm{main}}(x)$, $\Pi_t(x)$, $\pi_t(x)$, $\pi_t^*(x)$, $U_t(x)$, $C_t(x)$, $A_t(x)$ and so on.

2.3.6 Bad vertices

As mentioned in the proof overview, we will embed vertices of H into G one at a time, choosing each vertex in a way that allows for the future continuation of the embedding. This 'allowing for future continuation' comes in two parts, 'local' conditions which broadly say that we embed a vertex of H in a way which allows for the embedding of other vertices nearby in H, and 'global' conditions which say that the embedding of many (distant) vertices of H do not tend to cause problems for any given vertex. We

handle the latter probabilistically, and this constitutes the main work of this paper. The purpose of this subsection is to define precisely what we require locally, by specifying which vertices of G are bad for $x \in V(H)$, that is, not suitable for embedding x to, and to show that this is always a small set of vertices (in comparison to the set $A(x)$). In fact, the definition of badness also depends on a set $Q \subseteq V(H)$. In the random greedy algorithm (outlined in the proof overview, Section 1.4), Q will be the current queue, while at other times Q will simply be $V(H)$.

Definition 2.26 (Bad vertices, badness condition). Let ψ be a good partial embedding and $Q \subseteq V(H)$ be a set of unembedded vertices. The vertex v is called *bad for x with respect to ψ and Q* if the extension $\psi \cup \{x \to v\}$ is not a good partial embedding or there is an unembedded neighbour y of x with $y \notin Q$ such that

$$\deg_G\left(v; A^{\mathrm{main}}(y)\right) < (d - \varepsilon')p|A^{\mathrm{main}}(y)|. \tag{2.5}$$

We stress that in this equation we take $A^{\mathrm{main}}(y)$ with respect to the embedding ψ (and not the extended embedding). When ψ and Q are clear from the context (as they always will be) we let $B(x)$ be the set of vertices in $C(x)$ which are bad for x with respect to ψ and Q. When we have a sequence of embeddings $\psi_0, \ldots$ and a sequence of sets $Q_0, \ldots$ then we write, again, $B_t(x)$ for the set of bad vertices for x with respect to ψ_t and Q_t. We will also refer to (2.5) as the *badness condition*.

The following lemma provides control over the bad vertices with respect to a good partial embedding. We remark that the definition of the value D in this lemma is somewhat complex, because we want to apply this lemma in the proofs of all three of our blow-up lemmas, in particular the blow-up lemma for degenerate graphs.

Lemma 2.27 (Few bad vertices). *We assume the General Setup. Let ψ be a good partial embedding and let $Q \subseteq V(H)$ be a set of vertices such that for each unembedded $x \in V(H) \setminus Q$ we have*

$$|A^{\mathrm{main}}(x)| \geq \tfrac{1}{2}\mu(d - \varepsilon')^{\pi^*(x)}p^{\pi^*(x)}|V^{\mathrm{main}}(x)|.$$

Given an unembedded $x \in V(H)$, let D be such that for all unembedded $y, z \in V(H)$ we have

(i) $D \geq \pi^*(x) + 1$,

(ii) *If $xy \in E(H)$ then $D \geq \pi^*(y) + 1$,*

(iii) *If $xy, yz \in E(H)$ then $D \geq \pi^*(y) + 2, \pi^*(z) + 1$,*

(iv) *If $xy, yz, xz \in E(H)$ then $D \geq \pi^*(x) + 2, \pi^*(y) + 2, \pi^*(z) + 2$.*

Then the following holds.

(a) *If all neighbours of x are embedded then no vertex in $A(x)$ is bad for x with respect to ψ and Q.*

(b) *Suppose that Γ satisfies $\mathrm{NS}(\varepsilon, r_1, D)$ and $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$. Then at most $20\Delta^2\varepsilon'p^{\pi^*(x)}|V(x)|$ vertices of $C(x)$ are bad for x with respect to ψ and Q.*

Proof. We require

$$\mu \leq \tfrac{1}{6}, \qquad \varepsilon' \leq \frac{\mu d^\Delta \zeta}{1000\kappa 4^\Delta} \qquad \text{and} \qquad \varepsilon \leq \frac{d\varepsilon'}{\kappa}.$$

For (a), observe that all the conditions for a vertex v to be in $B(x)$ are trivially false.

For part (b), we need to consider all of the possible reasons v could be bad for x with respect to ψ and Q. It could be that there is some unembedded neighbour $y \in V(H) \setminus Q$ of x such that the badness condition (2.5) holds. The pair $\big(U(x), U(y)\big)$ is (ε', d, p)-regular in G by $(GPE\,4)$. By assumption and because $(GPE\,2)$ holds for y with respect to ψ we have

$$
\begin{aligned}
|A^{\mathrm{main}}(y)| &\geq \frac{1}{2}\mu(d-\varepsilon')^{\pi^*(y)}p^{\pi^*(y)}|V^{\mathrm{main}}(y)| = \frac{1}{2}\mu(1-3\mu)(d-\varepsilon')^{\pi^*(y)}p^{\pi^*(y)}|V(y)| \\
&\geq \frac{1}{2}\mu(1-3\mu)(d-\varepsilon')^{\pi^*(y)}p^{\pi^*(y)}\frac{|U(y)|}{(1+\varepsilon)^{\pi^*(y)}p^{\pi^*(y)}} \geq \varepsilon'|U(y)|.
\end{aligned}
$$

So the badness condition (2.5) holds for at most $\varepsilon'|U(x)| \leq 2\varepsilon' p^{\pi^*(x)}|V(x)|$ vertices v of $C(x)$, where we use $(GPE\,2)$ and $(1+\varepsilon)^{\Delta} \leq 2$. Since x has at most Δ unembedded neighbours, in total there are at most $2\Delta\varepsilon' p^{\pi^*(x)}|V(x)|$ vertices v of $C(x)$ such that the badness condition (2.5) holds for some unembedded neighbour of x.

It could be that $\psi \cup \{x \to v\}$ is not a good partial embedding. Since $v \in C(x) \subseteq I_x$, $(GPE\,1)$ cannot fail. We next consider $(GPE\,3)$. Since ψ is a good partial embedding, $(GPE\,3)$ cannot fail for any vertex y which is not a neighbour of x (because the candidate sets of these vertices do not change). So let y be any neighbour of x in H. Recall that $\big(U(x), U(y)\big)$ is (ε', d, p)-regular in G by $(GPE\,4)$. We have

$$
\begin{aligned}
|C(y)| &\overset{(GPE\,3)}{\geq} (1-\varepsilon')(dp-\varepsilon'p)^{\pi(y)}|I_y| \overset{(G\,5)}{\geq} (1-\varepsilon')(dp-\varepsilon'p)^{\pi^*(y)}\zeta|V(y)| \\
&\overset{(GPE\,2)}{\geq} (1-\varepsilon')(d-\varepsilon')^{\Delta}(1+\varepsilon')^{-\Delta}\zeta|U(y)| \geq \varepsilon'|U(y)|.
\end{aligned}
$$

Since $C(x) \subseteq U(x)$, there are at most $\varepsilon'|U(x)| \leq 2\varepsilon' p^{\pi^*(x)}|V_i|$ vertices $v \in C(x)$ such that $|C(y) \cap N_G(v)| < (d-\varepsilon')p|C(y)|$. Since $(d-\varepsilon')p|C(y)| \geq (1-\varepsilon')(dp-\varepsilon'p)^{\pi(y)+1}|I(y)|$, there are at most $2\varepsilon' p^{\pi^*(x)}|V(x)|$ vertices $v \in C(x)$ such that the first condition in $(GPE\,3)$ fails for y with respect to $\psi \cup \{x \to v\}$. We can analogously argue for each of the other four conditions. Since x has at most Δ unembedded neighbours, we conclude that for all but at most $10\Delta\varepsilon' p^{\pi^*(x)}|V_i|$ vertices $v \in C(x)$ we have that $(GPE\,3)$ is satisfied with respect to $\psi \cup \{x \to v\}$.

Now we turn to $(GPE\,2)$. Again, it suffices to consider neighbours y of x. As $(GPE\,2)$ holds with respect to ψ, we see that $|U(y)| \geq (p-\varepsilon p)^{\pi^*(y)}|V(y)| \geq \varepsilon p^{D-1}n/r_1$ since $|V(y)| \geq n/(\kappa r_1)$ and by assumption (ii). It follows that we can use $\mathrm{NS}(\varepsilon, r_1, D)$ and assumption (i) to conclude that all but at most $\varepsilon p^{D-1}n/r_1^2 \leq \varepsilon' p^{\pi^*(x)}|V(x)|$ vertices $v \in C(x)$ are such that $|U(y) \cap N_\Gamma(v)| = (1 \pm \varepsilon)p|U(y)| = (p \pm \varepsilon p)^{\pi^*(y)+1}|V(y)|$, again using $(GPE\,2)$ for y with respect to ψ. Again, we can argue analogously for the other four conditions, and so in total for all but at most $5\Delta\varepsilon' p^{\pi^*(x)}|V(x)|$ vertices $v \in C(x)$ we have that $(GPE\,2)$ is satisfied with respect to $\psi \cup \{x \to v\}$.

It remains to consider $(GPE\,4)$. Again, we only need to consider edges $yz \in E(H)$ between unembedded vertices such that $y \in N_H(x)$ or $z \in N_H(x)$, or both. Consider first the case that only $y \in N_H(x)$. The case that only $z \in N_H(x)$ follows analogously. Since $(GPE\,2)$ is satisfied with respect to ψ we have $|U(y)| \geq (p-\varepsilon p)^{\pi^*(y)}|V(y)|$ and $|U(z)| \geq (p-\varepsilon p)^{\pi^*(z)}|V(z)|$, which together with assumption (iii) implies $|U(y)| \geq (p-\varepsilon p)^{D-2}n/(\kappa r_1) \geq \varepsilon' p^{D-2}n/r_1^2$ and $|U(z)| \geq \varepsilon' p^{D-1}n/r_1^2$. Hence, since Γ satisfies $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$ and the pair $\big(U(y), U(z)\big)$ is $(\varepsilon_{\pi^*(y), \pi^*(z)}, d, p)$-regular in G because $(GPE\,4)$ is satisfied with respect to ψ, we get that there are at most $\varepsilon p^{D-1}n/r_1^2 \leq \varepsilon' p^{\pi^*(x)}|V(x)|$ vertices $v \in C(x)$

such that $\big(N_\Gamma(v)\cap U(y),U(z)\big)$ is not $(\varepsilon_{\pi^*(y)+1,\pi^*(z)},d,p)$-regular in G. Analogously, in the second case that $y,z\in N_H(x)$ we can use assumption *(iv)* to conclude that there are at most $\varepsilon p^{D-2}n/r_1^2 \leq \varepsilon' p^{\pi^*(x)}|V_i|$ vertices $v\in C(x)$ such that $\big(N_\Gamma(v)\cap U(y),N_\Gamma(v)\cap U(z)\big)$ is not $(\varepsilon_{\pi^*(y)+1,\pi^*(z)+1},d,p)$-regular in G. So, in total there are at most $\Delta^2\varepsilon' p^{\pi^*(x)}|V(x)|$ vertices $v\in C(x)$ such that *(GPE 4)* is not satisfied with respect to $\psi\cup\{x\to v\}$.

Summing up, there are at most

$$2\Delta\varepsilon' p^{\pi^*(x)}|V(x)| + 10\Delta\varepsilon' p^{\pi^*(x)}|V(x)| + 5\Delta\varepsilon' p^{\pi^*(x)}|V(x)| + \Delta^2\varepsilon' p^{\pi^*(x)}|V(x)|$$

$$\leq 20\Delta^2\varepsilon' p^{\pi^*(x)}|V(x)|$$

vertices $v\in C(x)$ such that v is bad for x with respect to ψ and Q. $\qquad\square$

2.4 The behaviour of random greedy algorithms

The proofs of our three blow-up lemmas each use, as crucial ingredient, a slightly different random greedy algorithm (RGA). In this section we collect four different lemmas which will be useful in conjunction with these RGAs.

Roughly speaking, each of our RGAs has the property that we embed vertices sequentially, at each time choosing an image of the next vertex uniformly at random in a not-too-small subset of its candidate vertices, and this, together with the General Setup (see Definition 2.23), is all we need to know about our RGAs for the purposes of this section. Nevertheless, the reader who wishes to see a concrete example of an RGA before reading the following lemmas should look at Algorithm 1 in Section 3.2, which is the simplest of our RGAs. The other two RGAs are Algorithm 3 in Section 4.2 and Algorithm 4 in Section 5.2.

Our first lemma (Lemma 2.28) constructs the vertex ordering which we shall use in the RGAs for the random graphs blow-up lemma (Lemma 1.21) and the blow-up lemma for bijumbled graphs (Lemma 1.25). The point of the ordering we use is that we need to analyse the way neighbours of (non-clique) buffer vertices are embedded during the RGA. Putting them first in the order, with each neighbourhood coming as an interval in the order, makes this possible. Even then, for Lemma 1.21 we need to insist that in degree-Δ buffers the final two neighbours of a buffer vertex are not adjacent, thus their embeddings do not affect each other much.

In the proof of our blow-up lemma for degenerate graphs (Lemma 1.23) we cannot apply this lemma, because we have to work with the order we are supplied with. Ultimately, the fact that we cannot expect these nice conditions from this supplied order is the reason why we cannot obtain stronger bounds on p in that lemma.

Lemma 2.28 (A good vertex order for an RGA). *Let H be a graph with $\Delta(H)\leq\Delta$ and $V(H)=X^{\text{buf}}\dot\cup X^{\text{main}}\dot\cup X^{\text{c}}$ such that $N(X^{\text{buf}})\subseteq X^{\text{main}}$ and each pair of vertices in X^{buf} has distance at least ten. Then there exists an ordering τ of X^{main} with the following properties.*

(a) For all $x\in X^{\text{main}}\setminus N(X^{\text{buf}})$ and $y\in N(X^{\text{buf}})$ we have $\tau(x)>\tau(y)$.

(b) For all $x\in X^{\text{buf}}$ we can enumerate $N(x)$ as $y_1,\ldots,y_b$ such that $\tau(y_{j+1})=\tau(y_j)+1$ for all $j\in[b-1]$. Moreover, if $\deg_H(x)=\Delta$ and x is not in a copy of $K_{\Delta+1}$ then $y_{\Delta-1}y_\Delta\notin E(H)$.

(c) *The neighbours of non-clique buffer vertices come before the neighbours of the clique-buffer vertices.*

Note that in this lemma we do not need to assume the General Setup, but we nevertheless use names for the parts of H corresponding to those from the General Setup.

Proof. We separate the vertices in X^{buf} into two classes, those not in copies of $K_{\Delta+1}$ and those in copies of $K_{\Delta+1}$. We take an enumeration $X^{\text{buf}} = \{x_1, \ldots, x_{|X^{\text{buf}}|}\}$ with the vertices not in copies of $K_{\Delta+1}$ coming first. We now create the ordering τ as follows. We start with the empty ordering. For each i in succession such that $x_i \in X^{\text{buf}}$ is not in a copy of $K_{\Delta+1}$, we append to τ the vertices $N_H(x_i)$ in some order such that if $\deg_H(x_i) = \Delta$ then the last two vertices in τ of $N_H(x_i)$ are not adjacent. Note that if $\deg_H(x) = \Delta$ then there is a pair of non-adjacent vertices in $N_H(x_i)$ because x_i is not in a copy of $K_{\Delta+1}$. Next, for each i in succession such that $x_i \in X^{\text{buf}}$ is in a copy of $K_{\Delta+1}$ we append to τ the vertices $N_H(x_i)$ in an arbitrary order. Finally, we append to τ any remaining vertices of X^{main} in an arbitrary order. $\qquad\square$

The next lemma states that provided we embed vertices uniformly at random into not-too-small sets, the candidate sets of unembedded vertices are likely to be distributed uniformly. In this lemma, as will be the case for much of the rest of the paper, we have a sequence $\psi_0, \ldots$ of good partial embeddings, and we want to talk about the various sets and quantities defined in Section 2.3.5, such as candidate sets, with respect to each of these good partial embeddings. As mentioned in Section 2.3.5, we will follow the convention that (for example) $C_t(x)$ is the candidate set of x with respect to ψ_t. Recall also, that the set $X_i^* \subseteq X_i$ denotes image restricted vertices in X_i (see Definition 1.20).

Lemma 2.29 (Uniform distribution of candidate sets)**.** *We assume the General Setup. Suppose that for some integer T we have a sequence $\psi_0, \psi_1, \ldots, \psi_T$ of good partial embeddings, where ψ_0 is the trivial partial embedding and each ψ_t is obtained from ψ_{t-1} by embedding some vertex x in $V(H) \setminus \text{Dom}(\psi_{t-1})$ to a uniform random vertex from a subset S of $C_{t-1}(x)$ with $|S| \geq \frac{1}{10}\mu\zeta(dp)^{\pi_{t-1}^*(x)}|V(x)|$.*

Then the following holds with probability at least $1 - 2^{-n/(\kappa r_1)}r_1$. For every $i \in [r]$ and every set $W \subseteq V_i$ of size at least $\rho|V_i|$, the number of vertices $x \in X_i \setminus X_i^*$ (i.e. vertices which are not image restricted) such that there exists $t = t(x)$ when x is unembedded and we have

$$|C_t(x) \cap W| < \left(dp - \varepsilon'p\right)^{\pi_t^*(x)}|W| \tag{2.6}$$

is at most $\rho|X_i|$.

This lemma corresponds to the Main Lemma of [56], and its proof is similar in spirit. The proof of this lemma exploits the fact that when a vertex y is embedded, condition (2.6) might become true for at most one vertex from $X_i \setminus X_i^*$ (since by (H2) the vertices in X_i have distance at least 10). Moreover, by the condition of always embedding into not too small subsets and by the regularity property (GPE4), the vertex y makes an unlucky choice with a very small probability, so that many such choices are extremely unlikely.

Proof of Lemma 2.29. We require

$$\varepsilon' < \min\left(4^{-\Delta}d^{\Delta}\rho, \frac{\mu\zeta d^{\Delta}}{20\Delta}2^{-4/\rho}\right) \quad \text{and} \quad 2^{-n/\kappa r_1}r_1 < 1.$$

Fix a set $W \subseteq V_i$ of size at least $\rho|V_i|$, and a set $X' \subseteq X_i \setminus X_i^*$ of size $\rho|X_i|$. We aim to show that the probability of the following event is at most $2^{-4|X_i|}$. For each $x \in X'$ there is a $t = t(x)$ such that x satisfies (2.6). The desired result then follows by taking the union bound over all i, subsets W of V_i and X' of X_i.

If x is any vertex of X', then we have $C_0(x) = V_i$ since x is not image restricted, thus $|C_0(x) \cap W| = |W|$ and hence (2.6) is false for x. If there is a $t(x)$ such that x satisfies (2.6), then we can fix t to be the smallest integer such that (2.6) is true for $t + 1$ and x. Since the candidate set of x changes only when a neighbour of x is embedded, it follows that the vertex y that is embedded to create ψ_{t+1} from ψ_t is a neighbour of x in H and thus $\pi_{t+1}^*(x) = \pi_t^*(x) + 1$. Moreover, since equation (2.6) becomes true for x, the vertex y is embedded to a vertex w such that

$$\deg_G\left(w; C_t(x) \cap W\right) < (d - \varepsilon')p\left|C_t(x) \cap W\right|, \tag{2.7}$$

as otherwise we would still have

$$|C_{t+1}(x) \cap W| \geq ((d - \varepsilon')p)|C_t(x) \cap W| \geq ((d - \varepsilon')p)^{\pi_{t+1}^*(x)}|W|.$$

Since ψ_t is a good partial embedding, by *(GPE 4)* the pair $\left(U_t(x), U_t(y)\right)$ is an (ε', d, p)-regular pair in G, and by *(GPE 2)* we have $\left|U_t(x)\right| = (p \pm \varepsilon' p)^{\pi_t^*(x)}|V_i|$, and $\left|U_t(y)\right| = (p \pm \varepsilon' p)^{\pi_t^*(y)}|V(y)|$. Since (2.6) is false for x at time t, we have $\left|C_t(x) \cap W\right| \geq ((d - \varepsilon')p)^{\pi_t^*(x)}|W|$, and so

$$\left|C_t(x) \cap W\right| \geq \left(\frac{d - \varepsilon'}{1 + \varepsilon'}\right)^{\pi_t^*(x)} \rho|U_t(x)| \geq \varepsilon'|U_t(x)|,$$

by the requirements on the constants at the beginning of the proof. We conclude by (ε', d, p)-regularity of $\left(U_t(x), U_t(y)\right)$ that at most $\varepsilon'\left|U_t(y)\right| \leq 2\varepsilon'p^{\pi_t^*(y)}|V(y)|$ vertices w of $U_t(y)$ satisfy (2.7).

Since ψ_{t+1} is created by embedding y uniformly at random into a subset of $C_t(y) \subseteq U_t(y)$ of size at least $\frac{1}{10}\mu\zeta(dp)^{\pi_t^*(y)}|V(y)|$, the probability of embedding y to a vertex w satisfying (2.7), conditioning on any history, is at most

$$\frac{2\varepsilon'p^{\pi_t^*(y)}|V(y)|}{\frac{1}{10}\mu\zeta(dp)^{\pi_t^*(y)}|V(y)|} \leq 20\varepsilon'\mu^{-1}\zeta^{-1}d^{-\Delta}. \tag{2.8}$$

Next we argue that the probability that for each $x \in X'$ there is a first time $t = t(x)$ such that x satisfies (2.6) is at most

$$\left(20\Delta\varepsilon'\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^{\rho|X_i|} \leq 2^{-4|X_i|}.$$

Let us denote this event by $\mathcal{E}_{X'}$. Observe that $\mathcal{E}_{X'}$ is split into $\Delta^{|X'|}$ events (since $\Delta(H) \leq \Delta$) by specifying for each $x \in X'$ a neighbour y_x of x whose embedding occurs at time $t(x)$. Let $(y_x)_{x \in X'}$ be any such assignment, and $\mathcal{E}_{X',(y_x)}$ be the corresponding event. We aim to bound $\mathbb{P}(\mathcal{E}_{X',(y_x)})$.

By *(H 2)* the vertices of X_i are at distance at least ten in H, so the vertices $(y_x)_{x \in X'}$ are distinct. The corresponding conditional probabilities thus multiply, and we have

$$\mathbb{P}(\mathcal{E}_{X',(y_x)}) \leq \left(20\varepsilon'\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^{|X'|}$$

by (2.8). Applying the union bound over the events $\mathcal{E}_{X',(y_x)}$ we conclude

$$\mathbb{P}(\mathcal{E}_{X'}) \leq \left(20\Delta\varepsilon'\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^{\rho|X_i|} \leq 2^{-4|X_i|}.$$

where the final inequality is by choice of ε'.

Taking the union bound over the at most $2^{|V_i|} = 2^{|X_i|}$ choices of W in V_i and $2^{|X_i|}$ choices of X' in X_i, we see that the probability that, for any fixed i, there exist subsets W of V_i and X' of X_i, of sizes at least $\rho|V_i|$ and $\rho|X_i|$ respectively, such that each vertex x of X' satisfies (2.6) at some time t, is at most $2^{2|X_i|}2^{-4|X_i|} = 2^{-2|X_i|}$. Now we have $|X_i| \geq n/(\kappa r_1)$, and $n \geq n_0$ where n_0 is chosen large enough such that $2^{-n_0/(\kappa r_1)}r_1 < 1$. Thus taking the union bound over the at most r_1 choices of i, we conclude that the probability that there exists i, and a subset W of V_i such that there are $\rho|X_i|$ vertices x of $X_i \setminus X_i^*$ each of which satisfies (2.6) at some time t, is at most $2^{-n/(\kappa r_1)}r_1$ as desired. $\qquad\square$

The next lemma, Lemma 2.30, shows that, again provided we embed vertices uniformly into not-too-small sets, we do not tend to cover vertex neighbourhoods in G disproportionately fast. Specifically, if $ij \in E(R')$ then, by (G3), each $v \in V_i$ has a large G-neighbourhood in V_j. At some time T when only a small but linear fraction of each part of H has been embedded, it is very likely that less than half of this G-neighbourhood is in the image $\text{Im}(\psi_T)$ of the current partial embedding. One should think of this as: early on in the embedding process, the minimum degree conditions (G3) provided by super-regularity are preserved. This fact is trivially satisfied in the dense blow-up lemma [56], but it is far from trivial in our case.

We will apply this lemma in the proofs of the random graphs blow-up lemma (Lemma 1.21) as well as the jumbled graphs blow-up lemma (Lemma 1.25). For the proof of the degenerate graphs blow-up lemma, on the other hand, we cannot use it; but we will replace it by a suitable modification of the ideas in its proof.

The idea behind Lemma 2.30 is as follows. As discussed in the proof overview (Section 1.4), we will need our RGAs to guarantee that each $v \in V_i$ is a candidate for many vertices $x \in X_i^{\text{buf}}$ in order to complete the embedding. This means we need it to be not too unlikely that $N_H(x)$ is embedded to $N_G(v)$ for any given $x \in X_i^{\text{buf}}$, and as a first step to showing this, it is necessary to show that we have not covered $N_G(v)$ with embedded vertices before we get around to embedding $N_H(x)$. In the proofs of the random graphs blow-up lemma and the bijumbled blow-up lemma, we show this by embedding $N(X^{\text{buf}})$ first and applying Lemma 2.30.

Lemma 2.30 (Preservation of super-regularity). *We assume the General Setup. Suppose that Γ has $\text{NS}(\varepsilon, r_1, B)$. Suppose that for some T we have a sequence $\psi_0, \psi_1, \ldots, \psi_T$ of good partial embeddings, where ψ_0 is the trivial partial embedding and each ψ_t is obtained from ψ_{t-1} by embedding one vertex $x_t \in V(H)$ to a uniform random vertex from a subset S of $C_{t-1}(x_t)$ with $|S| \geq \frac{1}{10}(dp)^{\pi_{t-1}^*(x_t)}|V(x_t)|$. Suppose furthermore that for each t the vertex x_t has at most $B-2$ neighbours in $\text{Dom}(\psi_{t-1})$, and that for each $i \in [r]$ we have $\left|\text{Dom}(\psi_T) \cap X_i\right| \leq 8\mu\kappa\Delta_{R'}|X_i|$.*

Then with probability at least $1 - \exp(-\varepsilon pn/r_1)$, for each $i \in [r]$, each vertex $v \in V_i$, and each j such that $ij \in E(R')$ we have $\left|N_G(v; V_j^{\text{main}}) \setminus \text{Im}(\psi_T)\right| \geq \frac{1}{2}\deg_G(v; V_j)$.

In the proof of the bijumbled blow-up lemma we set $B = \Delta + 1$ and will show that this is enough for Lemma 2.30 to handle all of $N(X^{\text{buf}})$. In the proof of the random graphs blow-up lemma, by contrast, we set $B = \Delta$, which turns out to be good enough to handle buffers of degree up to and including Δ which are not clique buffers. This is one of the reasons why we have to handle clique buffers differently.

The proof of Lemma 2.30 involves estimating the probability of embedding x_t to $N_G(v)$, conditioned on ψ_{t-1}. We show that either this probability is small, or that a previously embedded neighbour of x_t was

embedded oddly, which is guaranteed to be a low-probability event. In either case, for x_t to be embedded to $N_G(v)$ an unlikely event must occur. For $N_G(v)$ to be substantially filled up, many of these events, which are sequentially dependent, have to occur. The sequential dependence lemma, Lemma 2.2, then shows this is unlikely enough to take a union bound over all choices of v.

Proof. We require

$$\mu < \tfrac{d^\Delta}{1320\kappa\Delta_{R'}(\Delta+1)}, \quad \varepsilon < \tfrac{\mu d}{10\kappa} \quad \text{and} \quad r_1 n \exp(-\varepsilon pn/r_1) < 1.$$

Fix $ij \in E(R')$ and a vertex $v \in V_i$. We first estimate the conditional probability that $x_t \in X_j \cap \mathrm{Dom}(\psi_T)$ is embedded to $N_G(v) \cap V^{\mathrm{main}}(x_t)$, given the history. Because x_t is embedded uniformly at random to a subset of $C_{t-1}(x_t)$ of size at least $\frac{1}{10}(dp)^{\pi^*_{t-1}(x_t)}|V_j|$, this probability is at most

$$\frac{\left|U^{\mathrm{main}}_{t-1}(x_t) \cap N_G(v)\right|}{\frac{1}{10}(dp)^{\pi^*_{t-1}(x_t)}|V_j|}, \tag{2.9}$$

which we will show to be at most $20d^{-\Delta}\frac{\deg_G(v;V_j)}{|V_j|}$.

This clearly holds when $\left|U^{\mathrm{main}}_{t-1}(x_t) \cap N_G(v)\right| \le (p+\varepsilon p)^{\pi^*_{t-1}(x_t)} \deg_G(v;V_j)$ and we therefore assume the opposite. In this case the estimate on the conditional probability (2.9) that x_t occupies a vertex from $N_G(v)$ could be as great as 1, but we will show that this is an unlikely event. Specifically, for this to occur there must be a first neighbour y_t of x_t in H which is *embedded oddly* at time $t' < t$, that is, is such that $\left|U^{\mathrm{main}}_{t'-1}(x_t) \cap N_G(v)\right| \le (p+\varepsilon p)^{\pi^*_{t'-1}(x_t)} \deg_G(v;V_j)$ but (with $\pi^*_{t'}(x_t) = \pi^*_{t'-1}(x_t)+1$)

$$\deg_\Gamma\left(\psi_{t'}(y_t); U^{\mathrm{main}}_{t'-1}(x_t) \cap N_G(v)\right) > (p+\varepsilon p)^{\pi^*_{t'-1}(x_t)+1} \deg_G(v;V_j). \tag{2.10}$$

Let W be a superset of $U^{\mathrm{main}}_{t'-1}(x_t) \cap N_G(v)$ of size $(p+\varepsilon p)^{\pi^*_{t'-1}(x_t)} \deg_G(v;V_j) > \mu(d-\varepsilon)p^{B-1}|V_j|$, where the inequality uses the fact that $\pi^*_{t'-1}(x_t) \le B-2$ and (G3). By property $\mathrm{NS}(\varepsilon,r_1,B)$, which Γ satisfies, we see that there are at most $\varepsilon p^{B-1} n/r_1^2$ vertices Z of Γ which have this many neighbours in W and thus satisfy (2.10). Now $\pi^*_{t'-1}(y) \le B-2$ by the conditions of the lemma. We conclude that the probability of embedding y to a vertex of Z, conditioning on the history, is at most

$$\frac{\varepsilon p^{B-1} n/r_1^2}{\frac{1}{10}(dp)^{B-2}|V(y)|} \le 10\varepsilon d^{-\Delta}\kappa p < 10d^{-\Delta+1}p \le 20d^{-\Delta}\frac{\deg_G(v;V_j)}{|V_j|}, \tag{2.11}$$

where we use $|V(y)| \ge n/(\kappa r_1)$ in the first and the choice of ε in the second and (G3) in the third inequality.

We define a sequence of Bernoulli random variables $Y_1,\dots,Y_T$ as follows. Given the embedding ψ_{t-1}, if x_t is a neighbour of a vertex in $X_j \cap \mathrm{Dom}(\psi_T)$ none of whose previous neighbours were oddly embedded, and x_t is oddly embedded, we set $Y_t = 1$. If x_t is in X_j, none of its previous neighbours were oddly embedded, and x_t is embedded to $N_G(v;V_j)$, we set $Y_t = 1$. Otherwise we set $Y_t = 0$. By assumption, the total number of Y_t which are not deterministically zero (that is, are in $X_j \cap \mathrm{Dom}(\psi_T)$ or are neighbours of such a vertex) is at most $8\mu\kappa\Delta_{R'}(\Delta+1)|X_i|$. Observe furthermore that for all of the Y_t which are not deterministically zero, we have just shown in (2.11) that Y_t is one with probability at most $20d^{-\Delta}\frac{\deg_G(v;V_j^{\mathrm{main}})}{|V_j|}$ conditioning on the history up to, but not including, embedding x_t. This history

determines Y_{t-1}, so that the $Y_1, \ldots, Y_T$ are sequentially dependent. It follows that we can apply Lemma 2.2, with

$$x = 8\mu\kappa\Delta_{R'}(\Delta+1)|X_j| \cdot 20d^{-\Delta}\frac{\deg_G(v;V_j)}{|V_j|} = 160\mu\kappa d^{-\Delta}\Delta_{R'}(\Delta+1)\deg_G(v;V_j)$$

and $\delta = 1$, to show that

$$\mathbb{P}\big(Y_1 + \cdots + Y_T \geq 2x\big) \leq 2\exp\big(-\tfrac{x}{3}\big) < \exp\big(-\tfrac{2\varepsilon pn}{r_1}\big),$$

where the final inequality is by choice of ε and since $\deg_G(v;V_j) \geq (1-3\mu)(d-\varepsilon)p|V_j|$ by (G3). By choice of μ, we conclude that with probability at most $\exp(-2\varepsilon pn/r_1)$ we have $Y_1 + \cdots + Y_T \geq \frac{1}{4}\deg_G(v;V_j)$. Now observe that for any vertex to be embedded to $N_G(v;V_j^{\mathrm{main}})$ one of these variables $Y_1, \ldots, Y_T$ must be one, and since vertices of X_j are at distance at least ten in H (by (H2)), no Y_t can be responsible for two different vertices of X_j being embedded to $N_G(v;V_j^{\mathrm{main}})$. Thus $Y_1 + \cdots + Y_T$ is an upper bound for $\mathrm{Im}(\psi_T) \cap N_G(v;V_j^{\mathrm{main}})$. Taking a union bound over the choices of j and of v, and using (G3), we see that with probability at least $1 - r_1 n\exp(-2\varepsilon pn/r_1) > 1 - \exp(-\varepsilon pn/r_1)$ the statement of the lemma holds. $\qquad\square$

Our final lemma in this section complements the above lemma, showing that provided $v \in V_i$ does not have $N_G(v)$ more than half covered by $\mathrm{Im}(\psi)$ early in the embedding, it is reasonably likely that $N_H(x)$ is embedded to $N_G(v)$ for any given $x \in X_i^{\mathrm{buf}}$. As with the previous lemma, it contains a parameter B which will be either set to Δ (in the proof of the random graphs blow-up lemma) or to $\Delta + 1$ (in the proof of the bijumbled graphs blow-up lemma), and again in the former case the consequence is that we cannot use it to deal with clique buffers (that is, when $N_H(x)$ is a copy of K_Δ). For the degenerate graphs blow-up lemma we again need to replace this lemma by suitable similar ideas.

Lemma 2.31 (Buffer neighbour embedding). *We assume the General Setup. Let $B \in \{\Delta, \Delta+1\}$ and assume that Γ has $\mathrm{NS}(\varepsilon, r_1, B)$ and $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, B)$. Given $v \in V_i$ and $x \in X_i^{\mathrm{buf}}$, let $y_1, \ldots, y_b$ be the neighbours of x, and suppose that if $b = B$ then $y_{b-1}y_b$ is not an edge of H. Let ψ_0 be a good partial embedding in which no vertex at distance two or less from x is embedded, and let $Q_0 \subseteq X^{\mathrm{main}}$. Suppose further that $\big|N_G(v;V_j^{\mathrm{main}}) \setminus \mathrm{Im}(\psi_0)\big| \geq \frac{1}{2}\deg_G(v;V_j)$ for each j such that $ij \in R'$.*

Let $\psi_1, \ldots, \psi_b$ be good partial embeddings and $Q_1, \ldots, Q_b$ subsets of X^{main}, such that for each $t = 1, \ldots, b$ the embedding ψ_t is obtained from ψ_{t-1} by embedding y_t uniformly at random into its available candidates minus the vertices which are bad for y_t, that is, into $A_{t-1}^{\mathrm{main}}(y_t) \setminus B_{t-1}(y_t)$, and such that we have

$$Q_t = Q_{t-1} \cup \big\{z \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_t) \colon \big|A_t^{\mathrm{main}}(z)\big| < \tfrac{1}{2}\mu(d-\varepsilon')^{\pi_t^*(z)} p^{\pi_t^*(z)} |V^{\mathrm{main}}(z)|\big\}.$$

Then with probability at least $\big(d^\Delta p/10\big)^b$, we have $\psi_b\big(N_H(x)\big) \subseteq N_G(v)$.

The proof of this lemma is quite long, but much of it is 'bookkeeping'. Briefly, the idea is as follows. In ψ_0, no neighbour of any y_j is embedded and hence each y_j has candidate set $C_0(y_j) = V(y_j)$. We know that $N_G\big(v;V(y_1)\big)$ is not covered by $\mathrm{Im}(\psi_0)$, so if we choose a uniform random vertex of $A_0(y_1)$ the probability of choosing a member of $N_G(v)$ is at least $dp/4$. However, the random greedy algorithm does not choose a random vertex of $A_0(y_1)$, but rather of $A_0(y_1) \setminus B_0(y_1)$. Thus we have to show that $B_0(y_1)$

does not cover too much of $N_G(v)$, which we can do using properties $(G\,2)$ and $(G\,3)$ of the partition of G, and $\mathrm{NS}(\varepsilon,r_1,B)$ and $\mathrm{RI}(\varepsilon,(\varepsilon_{a,b}),\varepsilon',d,r_1,B)$ which we assume of Γ, in much the same way as the proof of Lemma 2.27. In addition we have to show that some extra properties are preserved which allow us to analyse the embedding of $y_2,\ldots,y_b$, which we can show is likely in a similar way. We will refer to a good partial embedding with these extra properties as a *buffer-friendly partial embedding*. We conclude that the probability of embedding y_j to $N_G(v)$ and maintaining a buffer-friendly partial embedding is at least $d^\Delta p/10$, conditioning on the history, for each j. The conditional probabilities multiply, giving the desired result.

Proof. We require

$$\mu \le \tfrac{1}{6}, \qquad \varepsilon < (\varepsilon')^2, \qquad \varepsilon' \le \frac{\mu d^\Delta \zeta}{1000\kappa\Delta^2} \qquad \text{and} \qquad (d-\varepsilon')^{10\Delta} > \tfrac{1}{2}d^{10\Delta}.$$

For each $0 \le t \le b$ and each y_j, $j \in [b]$, we define $\hat{A}_t^{\mathrm{main}}(y_j) := A_t^{\mathrm{main}}(y_j) \cap N_G(v)$ and for any vertex $x' \in X$ we set $\hat{U}_t(x') := U_t(x') \cap N_\Gamma(v)$.

Observe that since y_t is embedded to a vertex of $C_{t-1}(y_t) \setminus B_{t-1}(y_t)$ for each $t = 1,\ldots,b$ we automatically maintain the property that ψ_t is a good partial embedding for each t (see Definition 2.26 of bad vertices). We now formulate five additional conditions on ψ_t for $t = 0,\ldots,b$, which allow us to show a lower bound on the desired probability inductively. If ψ_t satisfies these conditions we say that it is a *buffer-friendly partial embedding*.

$(BPE\,1)$ We have that $\{\psi_t(y_1),\ldots,\psi_t(y_t)\} \subseteq N_G(v)$,

$(BPE\,2)$ $|\hat{U}_t(y_k)| = (p \pm \varepsilon' p)^{\pi_t^*(y_k)} \deg_\Gamma(v;V(y_k))$ for $k = t+1,\ldots,b$,

$(BPE\,3)$ $|\hat{A}_t^{\mathrm{main}}(y_k)| \ge \tfrac{1}{2}(dp - \varepsilon' p)^{\pi_t^*(y_k)} \deg_G(v;V(y_k))$ for $k = t+1,\ldots,b$,

$(BPE\,4)$ for unembedded y_k and y_ℓ with $y_k y_\ell \in E(H)$ we have $(\hat{U}_t(y_k),\hat{U}_t(y_\ell))$ is $(\varepsilon_{\pi_t^*(y_k),\pi_t^*(y_\ell)},d,p)$-regular in G,

$(BPE\,5)$ for unembedded y_k and z with $y_k z \in E(H)$ the pair $(\hat{U}_t(y_k),U_t(z))$ is $(\varepsilon_{\pi_t^*(y_k),\pi_t^*(z)},d,p)$-regular in G.

Observe that these conditions are all satisfied for ψ_0, i.e. ψ_0 is a buffer-friendly partial embedding. Indeed, $(BPE\,1)$ is vacuously satisfied. $(BPE\,2)$ holds by definition of $\hat{U}$ and since no neighbours of any y_k are embedded. $(BPE\,3)$ holds by the definition of $\hat{A}^{\mathrm{main}}$, the assumption on $\deg_G(v;V(x_k) \setminus \mathrm{Im}(\psi_0))$ of the lemma, and since no neighbours of any y_k are embedded. If y_k and y_ℓ are adjacent in H then xy_ky_ℓ is a triangle in H with $x \in X_i^{\mathrm{buf}} \subseteq \tilde{X}_i$, so by $(G\,2)$ we have $(BPE\,4)$. Finally, again since $x \in X_i^{\mathrm{buf}} \subseteq \tilde{X}_i$, by $(G\,2)$ we have $(BPE\,5)$.

For each $t = 0,\ldots,b-1$ we let P_t be the set of *unfriendly* vertices u in $\hat{A}_t^{\mathrm{main}}(y_{t+1})$, that is, the vertices that are such that if $\psi_{t+1}(y_{t+1}) = u$ then ψ_{t+1} is not a buffer-friendly partial embedding.

We now show that, for any $t = 1,\ldots,b$, given a buffer-friendly partial embedding ψ_{t-1}, the probability that ψ_t is a buffer-friendly partial embedding is at least $d^\Delta p/10$. This clearly yields, by multiplying the conditional probabilities, the lower bound of $(d^\Delta p/10)^b$ claimed in the statement of the lemma. For the analysis it will suffice to show that only a tiny proportion of vertices from $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ are in $B_{t-1}(y_t)$ or P_{t-1}.

So assume that ψ_{t-1} is a buffer-friendly partial embedding. We first give some lower bounds on set sizes. Since ψ_{t-1} is a good partial embedding we have for an unembedded vertex z that

$$|U_{t-1}^{\mathrm{q}}(z)| \geq \mu(p - \varepsilon p)^{\pi_{t-1}^*(z)}|V(z)| \geq \varepsilon' p^{\pi_{t-1}^*(z)} n/r_1 \tag{2.12}$$

where the first inequality is by $(GPE\,2)$ and the second because $|V(z)| \geq n/(\kappa r_1)$ and by choice of ε'. The same lower bound also holds for $U_{t-1}(z), C_{t-1}(z)$ et cetera since these sets are all by $(GPE\,2)$ and $(GPE\,3)$ at least as large. Since in addition ψ_{t-1} is a buffer-friendly partial embedding, we have for any y_k with $k \geq t$ that

$$|\hat{U}_{t-1}(y_k)| \geq |\hat{A}_{t-1}^{\mathrm{main}}(y_k)| \geq \tfrac{1}{2}(dp - \varepsilon' p)^{\pi_{t-1}^*(y_k)} \deg_G\big(v; V(y_k)\big) \geq \varepsilon' p^{\pi_{t-1}^*(y_k)+1} n/r_1 \tag{2.13}$$

where the second inequality is by $(BPE\,3)$ and the third is by $(G\,3)$ and choice of ε'. We also have for $k \geq t$, using $(BPE\,2)$, $(BPE\,3)$, $(GPE\,2)$, $(G\,3)$ and the choice of ε and ε' that

$$|\hat{A}_{t-1}^{\mathrm{main}}(y_k)| \geq \max\big(\tfrac{1}{4}d^\Delta p|U_{t-1}(y_k)|, \tfrac{1}{4}d^\Delta|\hat{U}_{t-1}(y_k)|\big) \geq \varepsilon'|\hat{U}_{t-1}(y_k)|. \tag{2.14}$$

We are next going to estimate the number of bad vertices in $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$.

Claim 2.32. $|\hat{A}_{t-1}^{\mathrm{main}}(y_t) \cap B_{t-1}(y_t)| \leq 30\Delta^2 d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$.

Proof. We need to estimate the number of vertices w from $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ which are bad with respect to ψ_{t-1} and Q_{t-1}. For that we will consider all of the possible reasons.

First assume that there is some unembedded neighbour $z \in V(H) \setminus Q_{t-1}$ of y_t such that the badness condition

$$\deg_G\big(w; A^{\mathrm{main}}(z)\big) < (d - \varepsilon')p|A^{\mathrm{main}}(z)| \tag{2.15}$$

holds for w and z. The pair $\big(\hat{U}_{t-1}(y_t), U_{t-1}(z)\big)$ is $\big(\varepsilon_{\pi_{t-1}^*(y_t), \pi_{t-1}^*(z)}, d, p\big)$-regular in G by $(BPE\,5)$. Because $z \notin Q_{t-1}$ we have that

$$\begin{aligned}
|A_{t-1}^{\mathrm{main}}(z)| &\geq \tfrac{1}{2}\mu(d - \varepsilon')^{\pi_{t-1}^*(z)} p^{\pi_{t-1}^*(z)}|V^{\mathrm{main}}(z)| \\
&\geq \tfrac{1}{2}\mu(1 - 3\mu)(d - \varepsilon')^{\pi_{t-1}^*(z)} p^{\pi_{t-1}^*(z)}|V(z)| \geq \varepsilon'|U_{t-1}(z)|,
\end{aligned}$$

where the first inequality is by definition of Q_{t-1} in the statement of the lemma, the second inequality is by $(GPE\,2)$ and the last by choice of ε' and by $(GPE\,2)$. Therefore the badness condition (2.15) holds for at most

$$\varepsilon_{\pi_{t-1}^*(y_t), \pi_{t-1}^*(z)}|\hat{U}_{t-1}(y_t)| \overset{(2.14)}{\leq} 4d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$$

vertices w of $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$. Since y_t may have at most Δ unembedded neighbours, in total there are at most $4\Delta d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$ vertices w of $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ such that the badness condition (2.15) holds for some unembedded neighbour of y_t.

Next we need to estimate the number of vertices w from $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ such that $\psi_{t-1} \cup \{y_t \to w\}$ is not a good partial embedding (i.e. does not satisfy properties $(GPE\,1)$–$(GPE\,4)$). First observe that since $w \in \hat{A}_{t-1}^{\mathrm{main}}(y_t) \subseteq C_{t-1}(y_t) \subseteq I_{y_t}$, $(GPE\,1)$ cannot fail.

Next we turn to $(GPE\,2)$. It is sufficient to consider unembedded neighbours z of y_t, so $\pi_{t-1}^*(z) \leq \Delta - 1$. As the bound in (2.12) also holds for $U_{t-1}(z)$ we have $|U_{t-1}(z)| \geq \varepsilon' p^{\Delta-1} n/r_1$. Therefore, by the

neighbourhood size property $\text{NS}(\varepsilon, r_1, B)$, all but at most $\varepsilon p^{B-1} n / r_1^2$ vertices w from $\hat{A}_{t-1}^{\text{main}}(y_t)$ are such that $|U_{t-1}(z) \cap N_\Gamma(w)| = (1 \pm \varepsilon)p|U_{t-1}(z)| = (p \pm \varepsilon p)^{\tau_{t-1}^*(z)+1}|V(z)|$, referring to *(GPE 2)* for z with respect to ψ_{t-1}. Further we have $\pi_{t-1}^*(y_t) \le B - 2$ since y_t has at most $y_1, \dots, y_{B-2}$ as embedded neighbours (if $b = B$ then by assumption $y_b y_{b-1}$ is not an edge of H), so since $\varepsilon < (\varepsilon')^2$ we have

$$\varepsilon p^{B-1} n / r_1^2 \overset{(2.13)}{\le} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|.$$

We can argue analogously for the other four conditions of *(GPE 2)* obtaining in total that for all but at most $5\Delta\varepsilon'|\hat{A}_{t-1}^{\text{main}}(y_t)|$ vertices w from $\hat{A}_{t-1}^{\text{main}}(y_t)$, *(GPE 2)* is satisfied with respect to $\psi_{t-1} \cup \{y_t \to w\}$.

Now we consider *(GPE 3)*. Let z be any unembedded neighbour of y_t. Because z is at distance at most 2 from $x \in X^{\text{buf}}$, by *(H 5)* z is not image restricted, so $I_z = V(z)$. Since *(GPE 2)* and *(GPE 3)* hold for ψ_{t-1} we have:

$$|C_{t-1}(z)| \ge (1 - \varepsilon')(dp - \varepsilon' p)^{\pi_{t-1}^*(z)}|V(z)| \ge \varepsilon'|U_{t-1}(z)|.$$

By *(BPE 5)*, $(\hat{U}_{t-1}(y_t), U_{t-1}(z))$ is $(\varepsilon_{\pi_{t-1}^*(y_t), \pi_{t-1}^*(z)}, d, p)$-regular in G. So, there are at most

$$\varepsilon_{\pi_{t-1}^*(y_t), \pi_{t-1}^*(z)}|\hat{U}_{t-1}(y_t)| \overset{(2.14)}{\le} 4d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\text{main}}(y_t)|$$

vertices w in $\hat{U}_{t-1}(y_t)$ such that $|C_{t-1}(z) \cap N_G(w)| < (d - \varepsilon')p|C_{t-1}(z)|$. Similarly we argue for each of the other four conditions. Thus, since y_t has at most Δ unembedded neighbours, in total for all but at most $20\Delta d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\text{main}}(y_t)|$ vertices $w \in \hat{A}_{t-1}^{\text{main}}(y_t)$ we have that *(GPE 3)* is satisfied with respect to $\psi_{t-1} \cup \{y_t \to w\}$.

It remains to consider *(GPE 4)*. Again, we only need to consider edges $zz' \in E(H)$ between unembedded vertices such that $z \in N_H(y_t)$. Consider first the case that $z' \notin N_H(y_t)$. In this case we have $\pi^*(z) \le \Delta - 2$, $\pi^*(z') \le \Delta - 1$ and $\pi_{t-1}^*(y_t) \le \Delta - 2$ (since both x and z are unembedded neighbours of y_t). Since Γ satisfies property $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, B)$, using (2.12) and the fact that by *(GPE 4)* the pair $(U_{t-1}(z), U_{t-1}(z'))$ is $(\varepsilon_{\pi_{t-1}^*(z), \pi_{t-1}^*(z')}, d, p)$-regular in G, we see that there are at most

$$\varepsilon p^{\Delta-1} n / r_1^2 \overset{(2.13)}{\le} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$$

vertices $w \in \hat{A}_{t-1}^{\text{main}}(y_t)$ such that the pair $(N_\Gamma(w; U_{t-1}(z)), U_{t-1}(z'))$ is not $(\varepsilon_{\pi_{t-1}^*(z)+1, \pi_{t-1}^*(z')}, d, p)$-regular in G.

In the case that $z, z' \in N_H(y_t)$ we have $\pi_{t-1}^*(z') \le \Delta - 2$. We also have $\pi_{t-1}^*(y_t) \le B - 3$. This requires a little explanation. Observe that y_t has unembedded neighbours z and z', so $\pi_{t-1}^*(y_t) \le \Delta - 2$, and if $B = \Delta + 1$ then we are done. If $B = \Delta$, then observe that the embedded neighbours of y_t are contained in $\{y_1, \dots, y_{t-1}\}$. It follows that if $t \le \Delta - 2$, we have $\pi_{t-1}^*(y_t) \le \Delta - 3$ as desired. If $t = \Delta$, then x, z and z' are distinct (since the only unembedded neighbour of x is y_t) and we again have $\pi_{t-1}^*(y_t) \le \Delta - 3$. It remains to consider the case $t = \Delta - 1$. Again we are done if x, z and z' are distinct. If, however, (without loss of generality) we have $x = z$, then z' is an unembedded neighbour of x which is not y_t; in other words, we have $z' = y_\Delta$. But we assumed that $y_{\Delta-1} y_\Delta$ is not an edge of H, contradicting the assumption $y_t z' \in H$. We have thus justified $\pi_{t-1}^*(y_t) \le B - 3$, so by $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, B)$ we conclude that there are at most

$$\varepsilon p^{B-2} n / r_1^2 \overset{(2.13)}{\le} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$$

vertices $w \in \hat{A}_{t-1}^{\mathrm{main}}(y_t)$ such that the pair $\big(N_\Gamma(w;U_{t-1}(z)),N_\Gamma(w;U_{t-1}(z'))\big)$ is not $(\varepsilon_{\pi_{t-1}^*(z)+1,\pi_{t-1}^*(z')+1},d,p)$-regular. So, in total there are at most $\Delta^2\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$ vertices $w \in \hat{A}_{t-1}^{\mathrm{main}}(y_t)$ such that *(GPE 4)* is not satisfied with respect to $\psi_{t-1} \cup \{y_t \to w\}$.

Summing up, we conclude that

$$|\hat{A}_{t-1}^{\mathrm{main}}(y_t) \cap B_{t-1}(y_t)| \le 4\Delta d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)| + 5\Delta\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)| + 20\Delta d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)| + \Delta^2\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$$
$$\le 30\Delta^2 d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$$

as required. □

Next we will estimate the number of unfriendly vertices in $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$.

Claim 2.33. $|P_{t-1}| \le 10\Delta^2 d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$.

Proof. Observe that by definition, since ψ_{t-1} is a buffer-friendly partial embedding, for each $w \in \hat{A}_{t-1}^{\mathrm{main}}(y_t)$ the embedding $\psi_{t-1} \cup \{y_t \to w\}$ has *(BPE 1)*. For $t = b$ the remaining properties *(BPE 2)*–*(BPE 5)* are trivial, so we from now on assume $t \le b-1$.

For *(BPE 2)*, let $k \in \{t+1,\ldots,b\}$ be such that $y_t y_k \in E(H)$. Observe that $\pi_{t-1}^*(y_t) \le t-1 \le b-2 \le B-2$. Therefore, by $\mathrm{NS}(\varepsilon,r_1,B)$, by (2.13) and (2.14), for all but at most $\varepsilon p^{B-1}n/r_1^2 \le \varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$ vertices w we have $|N_\Gamma(w;\hat{U}_{t-1}(y_k))| = (1 \pm \varepsilon')p|\hat{U}_{t-1}(y_k)|$.

To estimate the number of vertices w that do not preserve *(BPE 3)*, we only need to consider those y_k with $k \ge t+1$ such that $y_t y_k \in E(H)$. We use that $\big(\hat{U}_{t-1}(y_t),\hat{U}_{t-1}(y_k)\big)$ is (ε',d,p)-regular in G by *(BPE 4)*. Further, by (2.14) we have $|\hat{A}_{t-1}^{\mathrm{main}}(y_t)| \ge \varepsilon'|\hat{U}_{t-1}(y_t)|$, so there are at most

$$\varepsilon'|\hat{U}_{t-1}(y_t)| \overset{(2.14)}{\le} 4d^{-\Delta}\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$$

vertices w in $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ such that $\deg_G(w;\hat{A}_{t-1}^{\mathrm{main}}(y_k)) < (d-\varepsilon')p|\hat{A}_{t-1}^{\mathrm{main}}(y_k)|$.

Next we consider *(BPE 5)*. Let $k \in \{t+1,\ldots,b\}$, and $z \in V(H)$ be such that $y_k z \in E(H)$. By *(BPE 5)* the pair $\big(\hat{U}_{t-1}(y_k),U_{t-1}(z)\big)$ is $(\varepsilon_{\pi_{t-1}^*(y_k),\pi_{t-1}^*(z)},d,p)$-regular in G. There are three cases to consider: $y_t y_k \in E(H)$, $y_t z \in E(H)$, and both.

In the first case (i.e. $y_t y_k \in E(H)$, $y_t z \notin E(H)$), we have $t < B-1$ because $y_{B-1}y_B$ is not an edge of H, so we conclude $\pi_{t-1}^*(y_t),\pi_{t-1}^*(y_k) \le t-1 \le B-3$, and $\pi_{t-1}^*(z) \le \Delta-1$. By (2.12) and (2.13), and $\mathrm{RI}(\varepsilon,(\varepsilon_{a,b}),\varepsilon',d,r_1,B)$, we see that for all vertices w of $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ but at most

$$\varepsilon p^{B-1}n/r_1^2 \overset{(2.13)}{\le} \varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|,$$

the pair $\big(N_\Gamma(w;\hat{U}_{t-1}(y_k)),U_{t-1}(z)\big)$ is $(\varepsilon_{\pi_{t-1}^*(y_k)+1,\pi_{t-1}^*(z)},d,p)$-regular in G.

In the second case (i.e. $y_t y_k \notin E(H)$, $y_t z \in E(H)$), we have $\pi_{t-1}^*(y_t),\pi_{t-1}^*(y_k) \le t-1 \le B-2$, and $\pi_{t-1}^*(z) \le t-1 \le \Delta-2$ and similarly, by $\mathrm{RI}(\varepsilon,(\varepsilon_{a,b}),\varepsilon',d,r_1,B)$, there are not more than $\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$ vertices w of $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ for which the pair $\big(\hat{U}_{t-1}(y_k),N_\Gamma(w;U_{t-1}(z))\big)$ is not $(\varepsilon_{\pi_{t-1}^*(y_k),\pi_{t-1}^*(z)+1},d,p)$-regular in G.

In the final case (i.e. $y_t y_k \in E(H)$, $y_t z \in E(H)$), again we have $t < B-1$, hence $\pi_{t-1}^*(y_t),\pi_{t-1}^*(y_k) \le B-3$, and $\pi_{t-1}^*(z) \le \Delta-2$, and again all but at most $\varepsilon'|\hat{A}_{t-1}^{\mathrm{main}}(y_t)|$ vertices w of $\hat{A}_{t-1}^{\mathrm{main}}(y_t)$ are such that $\big(N_\Gamma(w;\hat{U}_{t-1}(y_k)),N_\Gamma(w;U_{t-1}(z))\big)$ is $(\varepsilon_{\pi_{t-1}^*(y_k)+1,\pi_{t-1}^*(z)+1},d,p)$-regular in G.

In total, we see that for all but at most $\Delta^2 \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$ vertices w of $\hat{A}_{t-1}^{\text{main}}(y_t)$, the partial embedding $\psi_{t-1} \cup \{y_t \to w\}$ has *(BPE 5)*.

Finally, we handle *(BPE 4)*. Let $k, \ell \in \{t+1, \ldots, b\}$ be such that $y_k y_\ell \in E(H)$. We have

$$\pi_{t-1}^*(y_t), \pi_{t-1}^*(y_k), \pi_{t-1}^*(y_\ell) \leq t - 1 \leq B - 3\,.$$

By *(BPE 4)*, the pair $\left(\hat{U}_{t-1}(y_k), \hat{U}_{t-1}(y_\ell)\right)$ is $(\varepsilon_{\pi_{t-1}^*(y_k), \pi_{t-1}^*(y_\ell)}, d, p)$-regular in G. Without loss of generality we may assume $y_t y_k \in E(H)$, and again there are two cases to consider depending on whether $y_t y_\ell \in E(H)$ or not. As before, using (2.13), by $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, B)$, at most

$$\varepsilon p^{B-2} n / r_1^2 \overset{(2.13)}{\leq} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$$

vertices $w \in \hat{A}_{t-1}^{\text{main}}(y_t)$ are such that $\left(N_\Gamma(w; \hat{U}_{t-1}(y_k)), N_\Gamma(w; \hat{U}_{t-1}(y_\ell))\right)$ is not $(\varepsilon_{\pi_{t-1}^*(y_k)+1, \pi_{t-1}^*(y_\ell)+1}, d, p)$-regular in G. The other case follows similarly, and we conclude that for all but at most $\Delta^2 \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$ vertices w of $\hat{A}_{t-1}^{\text{main}}(y_t)$, the partial embedding $\psi_{t-1} \cup \{y_t \to w\}$ has *(BPE 4)*.

In total we get $|P_{t-1}| \leq \Delta \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)| + 4\Delta d^{-\Delta} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)| + \Delta^2 \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)| + \Delta^2 \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$ as claimed. □

We conclude from Claim 2.32 and Claim 2.33 that the number of vertices $w \in \hat{A}_{t-1}^{\text{main}}(y_t)$ such that w is bad for y_t with respect to ψ_{t-1} and Q_{t-1} or $w \in P_{t-1}$ is at most $40\Delta^2 d^{-\Delta} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)|$.

Now we can estimate the probability that ψ_t is a buffer-friendly partial embedding, conditioning on the history and on ψ_{t-1} being a buffer-friendly partial embedding. This event occurs if y_t is embedded to a member of $\hat{A}_{t-1}^{\text{main}}(y_t) \setminus \left(B_{t-1}(y_t) \cup P_{t-1}\right)$. The number of such vertices is at least

$$\left|\hat{A}_{t-1}^{\text{main}}(y_t)\right| - 40\Delta^2 d^{-\Delta} \varepsilon' |\hat{A}_{t-1}^{\text{main}}(y_t)| \geq \tfrac{1}{2} |\hat{A}_{t-1}^{\text{main}}(y_t)| \overset{(2.13)}{\geq} \tfrac{1}{4} (dp - \varepsilon' p)^{\pi_{t-1}^*(y_t)+1} |V(y_t)|\,,$$

while y_t is embedded into a set of size at most $U_{t-1}(y_t)$, which by *(GPE 2)* has size at most $(p + \varepsilon p)^{\pi_{t-1}^*(y_t)} |V(y_t)|$. We see that the desired conditional probability is at least

$$\frac{\tfrac{1}{4}(dp - \varepsilon' p)^{\pi_{t-1}^*(y_t)+1} |V(y_t)|}{(p + \varepsilon p)^{\pi_{t-1}^*(y_t)} |V(y_t)|} \geq \frac{d^\Delta p}{10}\,.$$

The statement of the lemma follows since the conditional probabilities multiply. □

Chapter 3

Proof of the blow-up lemma for random graphs

In this chapter we provide the proof of Lemma 1.21. This proof relies on four main lemmas: The RGA Lemma (Lemma 3.1), the queue embedding lemma (Lemma 3.2), the buffer defect lemma (Lemma 3.3, and the buffer embedding lemma (Lemma 3.4). These lemmas and the actual proof of Lemma 1.21 are provided in Section 3.1. The RGA Lemma is then proved in Section 3.2, the queue embedding lemma in Section 3.3, and the buffer defect lemma in Section 3.4. The proof of Lemma 3.4 is short and thus given already in Section 3.1.

3.1 Main lemmas and the proof of the blow-up lemma

We divide the proof of Lemma 1.21 into four lemmas, which correspond to our four different embedding stages described in the proof overview in Section 1.4. At the end of this section we show how these four lemmas together with Lemma 2.22 (good partitions lemma) imply Lemma 1.21.

The first of our lemmas encapsulates the outcome after applying the random greedy algorithm (RGA) which tries to embed X^{main} into V^{main}. How this RGA operates is explained in the proof of this lemma in Section 3.2. The lemma claims the existence of a good partial embedding with certain deterministic properties which we require and which the RGA with high probability produces. Recall that the RGA constructs the queues X_i^{q}.

Lemma 3.1 (RGA lemma). *We assume the General Setup. Suppose that Γ has $\text{NS}(\varepsilon, r_1, \Delta)$ and $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$. Then there is a good partial embedding ψ_{RGA} of H into G with the following properties. For each i, let $X_i^{\text{q}} := X_i^{\text{main}} \setminus \text{Dom}(\psi_{\text{RGA}})$. Then the following hold for each i. Let b be such that X_i^{buf} is a degree-b buffer.*

(RGA 1) All neighbours of all buffer vertices are embedded by ψ_{RGA}.

(RGA 2) Every vertex in $X_i \cap \text{Dom}(\psi_{\text{RGA}})$ is embedded to V_i^{main} by ψ_{RGA}.

(RGA 3) We have $|X_i^{\text{q}}| \leq 2\rho |X_i|$.

(RGA 4) For every set $W \subseteq V_i$ of size at least $\rho |V_i|$, there are at most $\rho |X_i|$ vertices in X_i^{buf} with fewer than $(dp)^b |W|/2$ candidates in W.

$(RGA\,5)$ *If X_i^{buf} is not a clique buffer, then every vertex in V_i is a candidate for at least $\mu\left(d^\Delta p/100\right)^b |X_i|$ vertices of X_i^{buf}.*

We shall next extend the good partial embedding provided by the RGA lemma and embed the queue vertices into the sets V_i^{q}. That this is possible is stated by the following lemma. As mentioned in the proof overview, the proof of this lemma relies on Γ not having a dense spot, which the congestion condition $\mathrm{CON}(\rho, r_1, \Delta)$ guarantees.

Lemma 3.2 (Queue embedding lemma). *We assume the General Setup. Suppose that Γ has $\mathrm{NS}(\varepsilon, r_1, \Delta)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$ and $\mathrm{CON}(\rho, r_1, \Delta)$. Let ψ be a good partial embedding whose image is disjoint from the sets V_i^{q}, and suppose that for each i we have a set $X_i^{\mathrm{q}} \subseteq X_i$ of size at most $2\rho |X_i|$. Then there is a good partial embedding ψ_{q} extending ψ such that*

$$\mathrm{Dom}(\psi_{\mathrm{q}}) \setminus \mathrm{Dom}(\psi) = X^{\mathrm{q}} \text{ and } \mathrm{Im}(\psi_{\mathrm{q}}) \setminus \mathrm{Im}(\psi) \subseteq V^{\mathrm{q}}.$$

Note that in the RGA lemma, Lemma 3.1, we do not have that all vertices v in V_i are candidates for many vertices of X_i^{buf} if X_i^{buf} is a clique buffer (and in fact our RGA may fail to produce this property). In this case we say that v has a *buffer defect* for X_i^{buf}. It follows from $(RGA\,4)$, however, that only at most $\rho |V_i|$ vertices of V_i can have a buffer defect for X_i^{buf}. Our next lemma, the buffer defect lemma, deals with fixing these buffer defects.

In order to apply this lemma we need to alter the embedding strategy described so far slightly. Indeed, before we begin the embedding, we will reserve some copies of $K_{\Delta+1}$ in H, each of which has a vertex in $\tilde{X}_i$. We call these the *reserved cliques* and denote the family of reserved cliques for $\tilde{X}_i$ by $\mathcal{K}_i$. We put the vertices of all reserved cliques into the set X^{c}, which is disjoint from X^{main} and X^{buf}. Then we apply the RGA lemma and the queue embedding lemma to obtain a good partial embedding of all of H but the reserved cliques X^{c} and the buffer vertices X^{buf}. Next, we use the buffer defect lemma to embed the reserved cliques (and, for technical reasons, also some vertices of X_i^{buf}) on vertices with buffer defects, to obtain a good partial embedding ψ_{good} in which only the vertices of X^{buf} remain unembedded and without any buffer defects left. Note that in the following lemma $A^{\mathrm{buf}}(x) = A(x) \cap V^{\mathrm{buf}}(x) = C(x) \cap V^{\mathrm{buf}}(x) \setminus \mathrm{Im}(\psi)$.

Lemma 3.3 (Buffer defect lemma). *We assume the General Setup. Suppose that Γ has $\mathrm{NS}(\varepsilon, r_1, \Delta)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$ and $\mathrm{CON}(\rho, r_1, \Delta)$. Let ψ be a good partial embedding whose image is disjoint from the sets V_i^{c} and whose domain contains $N(X^{\mathrm{buf}})$. Suppose that for each i we are given a family $\mathcal{K}_i$ of reserved cliques $K_{\Delta+1}$, not embedded by ψ, such that for each $K \in \mathcal{K}_i$ and $x, y \in K$ we have $x \in X_j$ and $y \in X_k$ with $jk \in R'$, and that either*

(a) *X_i^{buf} is a degree-b non-clique buffer with $\mathcal{K}_i = \emptyset$, and all vertices in V_i are candidate for at least $\mu\left(d^\Delta p/100\right)^b |X_i|$ vertices in X_i^{buf}, or*

(b) *X_i^{buf} is a clique buffer with $|\mathcal{K}_i| = 2\rho |X_i|$, and all but at most $\rho |V_i|$ vertices of V_i are candidate for at least $\mu(dp)^\Delta |X_i|/2$ vertices in X_i^{buf}.*

Let X_i^{c} be the set of vertices of X_i contained in some reserved clique from $\mathcal{K}_i$ and assume that $|X_i^{\mathrm{c}}| \leq 2\kappa(\Delta_{R'} + 1)\rho |X_i|$.

Then there is a good partial embedding ψ_{good} extending ψ such that for each i the following hold. Let b be such that X_i^{buf} is a degree-b buffer.

(*BD* 1) $X_i^c \subseteq \text{Dom}(\psi_{\text{good}})$.

(*BD* 2) *For each* $x \in X_i^{\text{buf}} \setminus \text{Dom}(\psi_{\text{good}})$ *we have* $|A^{\text{buf}}(x)| \geq \mu (dp)^b |V_i|/4$.

(*BD* 3) *Each* $v \in V_i \setminus \text{Im}(\psi_{\text{good}})$ *is a candidate for at least* $\mu (d^\Delta p/100)^b |X_i|$ *vertices in* $X_i^{\text{buf}} \setminus \text{Dom}(\psi_{\text{good}})$.

It remains to embed the buffer vertices, which is performed with the help of the following lemma.

Lemma 3.4 (Buffer embedding lemma)**.** *We assume the General Setup. Suppose that* Γ *has* $\text{CON}(\rho, r_1, \Delta)$. *Suppose that* ψ *is a good partial embedding, in which all the vertices* $N(X^{\text{buf}})$ *are embedded, and that for each* i *we have a subset* X_i' *of* X_i^{buf}, *and a subset* V_i' *of* $V_i \setminus \text{Im}(\psi)$ *with* $|X_i'| = |V_i'|$. *Suppose* X_i^{buf} *is a degree-b buffer, and for some* $\delta > 0$ *and* $\rho \leq \frac{\delta \mu d^b}{100\kappa}$ *we have*

(*BUF* 1) *for each* $x \in X_i'$ *we have* $|C(x) \cap V_i'| \geq \mu (dp)^b |V_i|/4$,

(*BUF* 2) *for each* $W \subseteq V_i'$ *with* $|W| \geq \rho|V_i|$, *there are at most* $\rho|X_i|$ *vertices in* X_i' *which do not have a candidate in* W, *and*

(*BUF* 3) *each* $v \in V_i'$ *is a candidate for at least* $\delta p^b |X_i|$ *vertices in* X_i'.

Then there is a good partial embedding ψ' *extending* ψ *such that* $\text{Dom}(\psi') = \text{Dom}(\psi) \cup \bigcup_i X_i'$ *and* $\text{Im}(\psi') = \text{Im}(\psi) \cup \bigcup_i V_i'$.

The proof of this lemma mainly consists of a straightforward check of Hall's condition and we provide it right away. We remark that we will use this lemma again in the proof of the degenerate graphs blow-up lemma, which is why we do not give an explicit constant in place of δ in (*BUF* 3).

Proof of Lemma 3.4. Since X^{buf} is independent in H (and the sets V_i' are disjoint in G), we can embed the vertices of X_i' for each i in succession, without affecting the vertices in any other X_i'. So we fix i and will show that X_i' can be embedded into V_i'.

Let $Y \subseteq X_i'$ be non-empty, and let W be the set of vertices in V_i' which are a candidate for some member of Y. We wish to verify Hall's condition, i.e. show that $|W| \geq |Y|$. We separate three cases.

First, suppose $0 < |Y| \leq \rho|X_i|$. If $|W| < |Y|$, then we can take a subset Y' of Y of size $|W|$. Recall that the candidate graph between Y' and W is the bipartite graph with edges yw for $y \in Y'$ and $w \in W$ with $w \in C(y)$, and that this is a subgraph of the underlying restriction graph between Y' and W with edges yw with $w \in U(y)$. Observe, moreover, that this underlying restriction graph is isomorphic to the congestion graph $\text{CG}(\Gamma, W, \mathcal{F})$ with $\mathcal{F} = \{\psi(N_H(y)) : y \in Y'\}$, which has edges wF with $w \in N_\Gamma^*(F)$. Now, since by (*BUF* 1) each vertex in Y' has at least $\mu (dp)^b |V_i|/4$ candidates in V_i', which must lie in W, the number of edges in the candidate graph between Y' and W is at least

$$\tfrac{1}{4}\mu (dp)^b |V_i||Y'| = \tfrac{1}{8}\mu (dp)^b |V_i||Y'| + \tfrac{1}{8}\mu (dp)^b |V_i||Y'|$$

$$> \frac{\mu d^b}{8\rho} p^b |Y'||W| + \frac{\mu d^b}{8\kappa r_1} p^b n |Y'|$$

$$> 7p^b |Y'||W| + \rho p^b n |Y'|/r_1 \,,$$

where the second line comes from substituting $|V_i| = |X_i| > |Y'|/\rho = |W|/\rho$ and using $|V_i| \geq n/(\kappa r_1)$, and the third follows from the bound on ρ. Hence $e(\text{CG}(\Gamma, W, \mathcal{F})) > 7p^b |Y'||W| + \rho p^b n |Y'|/r_1$, in contradiction to $\text{CON}(\rho, r_1, \Delta)$. We conclude that $|W| \geq |Y|$ in this case.

Second, if $\rho|X_i| < |Y| \le |X_i'| - \rho|X_i| = |V_i'| - \rho|V_i|$ and $|W| < |Y|$, then $|V_i' \setminus W| > \rho|V_i|$, so by *(BUF 2)* there are at most $\rho|X_i|$ vertices of X_i' which do not have candidates in $V_i' \setminus W$. In particular there is a vertex of Y with candidates in $V_i' \setminus W$, in contradiction to the definition of W. We conclude that $|W| \ge |Y|$ in this case as well.

Finally, suppose $|Y| > |X_i'| - \rho|X_i| = |V_i'| - \rho|V_i|$. The vertices $V_i' \setminus W$ are candidates only for vertices in $X_i' \setminus Y$, and each vertex in $V_i' \setminus W$ is a candidate for at least $\delta p^b |X_i|$ vertices in $X_i' \setminus Y$ by *(BUF 3)*. If $|W| < |Y|$, then $|V_i' \setminus W| > |X_i' \setminus Y|$, and we can take a set $\tilde{W} \subseteq V_i' \setminus W$ of size $|X_i' \setminus Y|$. Now the number of edges in the candidate graph between $\tilde{W}$ and $X_i' \setminus Y$ is at least

$$\delta p^b |X_i||\tilde{W}| > \tfrac{1}{2}\delta p^b\left(\tfrac{1}{\rho}|X_i' \setminus Y||\tilde{W}| + \tfrac{n}{\kappa r_1}|X_i' \setminus Y|\right)$$
$$> 7p^b|X_i' \setminus Y||\tilde{W}| + \rho p^b n|X_i' \setminus Y|/r_1$$

by essentially the same calculation as in the first case, using $|X_i| = |V_i| \ge n/(\kappa r_1)$. Taking $\mathcal{F}' = \left\{\psi\big(N_H(x)\big) \colon x \in X_i' \setminus Y\right\}$, the number of edges in the corresponding congestion graph $\mathrm{CG}(\Gamma, \tilde{W}, \mathcal{F}')$ is also at least this quantity, in contradiction to $\mathrm{CON}(\rho, r_1, \Delta)$.

This completes the verification of Hall's condition, so there is a partial embedding ψ' extending ψ with $\mathrm{Dom}(\psi') = \mathrm{Dom}(\psi) \cup X_i'$ in which the vertices X_i' are embedded to V_i'. Since X^{buf} is independent and all the vertices $N(X^{\mathrm{buf}})$ are embedded in ψ, it is trivially the case that ψ' is a good partial embedding. $\square$

We are now ready to prove Lemma 1.21. Summarising, the proof will go as follows. We will apply Lemma 2.22 to find partitions of G and H, and graphs R and R', satisfying the General Setup. We choose, for each i such that X_i^{buf} is a clique buffer a set of reserved cliques with vertices X^{c}. We can then apply the RGA lemma, the queue embedding lemma, the buffer defect lemma, and the buffer embedding lemma, in this order.

Proof of Lemma 1.21. First we choose constants as follows. Given Δ, $\Delta_{R'}^{\mathrm{BL}}$, Δ_J integers, α^{BL}, ζ^{BL} and $d > 0$, and $\kappa^{\mathrm{BL}} > 1$, we set $\vartheta = 0$, $\Delta_{R'} = 8(\Delta + \Delta_J)^{10}\Delta_{R'}^{\mathrm{BL}}$, $\alpha = \tfrac{1}{2}\alpha^{\mathrm{BL}}$, $\zeta = \tfrac{1}{2}\zeta^{\mathrm{BL}}$ and $\kappa = 2\kappa^{\mathrm{BL}}$. We now choose μ, ρ and $\varepsilon' > 0$ satisfying the conditions in Lemmas 2.22, 2.27, 3.1, 3.2, 3.3 and 3.4. For convenience we provide here sufficient choices:

$$\mu < \frac{d^\Delta}{1320\kappa\Delta_{R'}}, \text{ and } \rho \le \frac{\mu^2 d^{\Delta^2+1}}{250^{\Delta+1}\kappa\Delta_{R'}}, \text{ and } \varepsilon' \le \frac{\mu\zeta\rho d^\Delta}{32^{\Delta+2}2^{4/\rho}\kappa\Delta^2\Delta_{R'}}.$$

Now for input Δ, d and ε', Lemma 2.17 returns constants $\varepsilon_{a,b}$ and $\varepsilon > 0$. Here we additionally require that $\varepsilon < (\varepsilon')^2$. We let $\varepsilon^{\mathrm{BL}} = \tfrac{1}{16}(\Delta + \Delta_J)^{-10}\varepsilon$ and $\rho^{\mathrm{BL}} = \tfrac{1}{16}(\Delta + \Delta_J)^{-10}\rho$. Now Lemma 1.21 returns $\varepsilon^{\mathrm{BL}}$ and ρ^{BL}. Given r_1^{BL} we let $r_1 = 8(\Delta + \Delta_J)^{10}r_1^{\mathrm{BL}}$. We choose C sufficiently large for Lemma 2.17 with input Δ, d, ε', $T = r_1$ and ρ.

Given $p \ge C\left(\frac{\log n}{n}\right)^{1/\Delta}$, Lemma 2.17 states that a.a.s. $\Gamma = G(n,p)$ has properties

$$\mathrm{NS}(\varepsilon, r_1, \Delta), \mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta) \text{ and } \mathrm{CON}(\rho, r_1, \Delta),$$

respectively. From now on we will assume Γ is an n-vertex graph which satisfies these three properties.

Given a graph R_{BL} on $r_{\mathrm{BL}} \le r_1^{\mathrm{BL}}$ vertices, and a spanning subgraph R_{BL}' with $\Delta(R_{\mathrm{BL}}') \le \Delta_{R'}^{\mathrm{BL}}$, and graphs H and $G \subseteq \Gamma$ with vertex partitions $\mathcal{X}^{\mathrm{BL}}$ and $\mathcal{V}^{\mathrm{BL}}$, families of image restrictions $\mathcal{J}^{\mathrm{BL}}$ and of

image restricting vertices $\mathfrak{J}$, and a family of potential buffer vertices $\tilde{\mathcal{X}}^{\mathrm{BL}}$, suppose that the conditions of Lemma 1.21 are satisfied. Then Lemma 2.22 gives (with Δ in place of b) a graph R on $r \leq r_1$ vertices, a spanning subgraph R' with $\Delta(R') \leq \Delta_{R'}$, and κ-balanced size-compatible partitions $\mathcal{X}$ and $\mathcal{V}$ of H and G respectively, each part having size at least $n/(\kappa r_1)$, together with a family $\tilde{\mathcal{X}}$ of potential buffer vertices and $\mathfrak{J}$ of image restrictions, subsets X_i^{buf} of $\tilde{X}_i$ for each $i \in [r]$, and partitions $V_i = V_i^{\mathrm{main}} \dot{\cup} V_i^{\mathrm{q}} \dot{\cup} V_i^{\mathrm{c}} \dot{\cup} V_i^{\mathrm{buf}}$ for each $i \in [r]$ which satisfy the General Setup.

Now for each $i \in [r]$ such that X_i^{buf} is a clique buffer, in succession, we create a collection of reserved cliques $\mathcal{K}_i$ and a set X^{c} as follows. We choose $2\rho|X_i|$ vertices from $\tilde{X}_i$ which are in copies of $K_{\Delta+1}$ that do not contain image restricted vertices, vertices of X^{buf}, or vertices previously added to X^{c}. We add these copies of $K_{\Delta+1}$ to $\mathcal{K}_i$, and their vertices to X^{c}. Observe that this is possible since $(H\,8)$ guarantees that each $\tilde{X}_i$ contains at least $\frac{1}{2\Delta+4}\alpha|X_i|$ vertices in copies of $K_{\Delta+1}$ whose vertices are neither image restricted nor in X^{buf} and, moreover, by $(H\,1)$ the edges of these cliques lie along R', so the number of these cliques which are chosen for some $\mathcal{K}_j$ with $j < i$ is at most $\Delta_{R'} 2\rho\kappa|X_i| < \frac{1}{4\Delta+8}\alpha|X_i|$ by choice of ρ. We conclude the following properties for these reserved cliques.

(*RSC* 1) The sets $\{\mathcal{K}_i\}_{i \in [r]}$ are pairwise disjoint, and if X_i^{buf} is a clique buffer, then $\mathcal{K}_i$ contains $2\rho|X_i|$ cliques, each with one vertex in X_i, otherwise it is empty.

(*RSC* 2) For each $i \in [r]$, $K \in \mathcal{K}_i$ and $x, y \in K$ we have $x \in X_j$ and $y \in X_k$ where $jk \in R'$.

(*RSC* 3) For each i, the set X_i^{c} of vertices $x \in X_i$ contained in some reserved clique satisfies $|X_i^{\mathrm{c}}| \leq 2\kappa(\Delta_{R'}+1)\rho|X_i|$.

Notice that (*RSC* 2) holds by $(H\,1)$, since the first and second neighbours of vertices from $\tilde{X}_i$ go along the edges of R' by the definition of an (α, R')-buffer.

We let $X^{\mathrm{main}} = V(H) \setminus (X^{\mathrm{buf}} \cup X^{\mathrm{c}})$. We now begin the embedding of H into G. By Lemma 3.1, there is a good partial embedding ψ_{RGA} with properties (*RGA* 1)–(*RGA* 5). Letting $X^{\mathrm{q}} = X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_{\mathrm{RGA}})$, by (*RGA* 3) we have $|X_i^{\mathrm{q}}| \leq 2\rho|X_i|$ for each i, and by (*RGA* 2) we see that $\mathrm{Im}(\psi_{\mathrm{RGA}})$ is disjoint from each set V_i^{q}, so the conditions of Lemma 3.2 are met. Feeding ψ_{RGA} into Lemma 3.2 we obtain a good partial embedding ψ_{q} extending ψ_{RGA} whose domain is X^{main} and whose image is contained in $V^{\mathrm{main}} \cup V^{\mathrm{q}}$. By (*RGA* 1) all neighbours of all buffer vertices are in $\mathrm{Dom}(\psi_{\mathrm{RGA}})$, and therefore the candidate sets of all buffer vertices are the same with respect to ψ_{RGA} as with respect to ψ_{q}. In particular ψ_{q} satisfies (*RGA* 4) and (*RGA* 5).

We now verify the conditions for Lemma 3.3. Since ψ_{q} extends ψ_{RGA} and the vertices $\mathrm{Im}(\psi_{\mathrm{q}}) \setminus \mathrm{Im}(\psi_{\mathrm{RGA}})$ are embedded in the sets V_i^{q} it follows by (*RGA* 2) that $\mathrm{Im}(\psi_{\mathrm{q}})$ is disjoint from the sets $\bigcup_i V_i^{\mathrm{c}}$ and $\bigcup_i V_i^{\mathrm{buf}}$. By (*RGA* 5), if X_i^{buf} is a degree-b buffer for some b but is not a clique buffer, then all vertices in V_i are candidate for at least $\mu(d^{\Delta}p/100)^b|X_i|$ vertices in X_i^{buf} and $\mathcal{K}_i$ is empty by (*RSC* 1).

If on the other hand X_i^{buf} is a clique buffer, then $|\mathcal{K}_i| = 2\rho|X_i|$. Now let W be the set of vertices in V_i which are candidates for fewer than $\mu(dp)^{\Delta}|X_i|/2$ vertices of X_i^{buf}. We will show that $|W| < \rho|V_i|$. If $|W| \geq \rho|V_i|$, then by (*RGA* 4) there are at most $\rho|X_i|$ vertices in X_i^{buf} with fewer than $(dp)^{\Delta}|W|/2$ candidates in W. It follows that the average over $v \in W$ of the number of vertices of X_i^{buf} for which v is a candidate, is at least

$$\frac{1}{|W|}\left(|X_i^{\mathrm{buf}}| - \rho|X_i|\right)\frac{(dp)^{\Delta}|W|}{2} > \frac{1}{2}\mu(dp)^{\Delta}|X_i|,$$

where we used the facts $|X_i^{\mathrm{buf}}| = 4\mu|X_i|$ and $\rho < \mu$ in the inequality. Since at least one $v \in W$ attains at

least the average, we have a contradiction to the definition of W.

The conditions of Lemma 3.3 are thus satisfied, and by applying it we obtain a good partial embedding ψ_{good}. Since ψ_{good} extends ψ_q and all sets X_i^c are contained in $\text{Dom}(\psi_{\text{good}})$ by (*BD* 1), we conclude that the only vertices remaining unembedded are in X^{buf}.

Finally, for $i \in [r]$, let $X_i' = X_i^{\text{buf}} \setminus \text{Dom}(\psi_{\text{good}})$ and let X_i^{buf} be a degree-b buffer for some b. Let $V_i' = V_i \setminus \text{Im}(\psi_{\text{good}})$. Because $|X_i| = |V_i|$ we have $|X_i'| = |V_i'|$. Because ψ_{good} satisfies (*BD* 2), we have (*BUF* 1). Because ψ_{good} satisfies (*RGA* 4), in particular we have (*BUF* 2). Finally, because ψ_{good} satisfies (*BD* 3) we have (*BUF* 3) with $\delta = \mu(d^\Delta/100)^b$. Thus by Lemma 3.4 there is an embedding ψ' extending ψ_{good} which embeds H into G. $\qquad\square$

3.2 Proof of the random graphs RGA lemma

In this section we describe the random greedy algorithm (RGA) and prove that it produces a partial embedding which satisfies the assertions of the RGA lemma, Lemma 3.1, with high probability. This is Algorithm 1. It embeds vertices x of H sequentially, following an order τ given by Lemma 2.28 (a good vertex order for an RGA). In doing so, it builds up a sequence $\psi_0, \dots$ of good partial embeddings, and a *queue* of vertices which it will not embed; we let Q_t be the queue at time t (i.e. corresponding to ψ_t). Recall that by $A_t^{\text{main}}(x)$ we mean the set $A^{\text{main}}(x)$ with reference to the partial embedding ψ_t. We let $B_t(x)$ denote the set of bad vertices (defined in Section 2.3.6) with respect to ψ_t and Q_t. As mentioned in the proof overview (Section 1.4), to create ψ_{t+1} from ψ_t we embed some vertex x uniformly at random into the set $A_t^{\text{main}}(x) \setminus B_t(x)$, the set of available candidate vertices in V^{main} minus the bad vertices; and we add y to the queue if the set $A_t^{\text{main}}(y) \setminus B_t(y)$ gets small.

Note that exactly t vertices are embedded in ψ_t, though these vertices are not necessarily the first t vertices of τ because vertices in Q_t are skipped. The queue set $Q_{t_{\text{RGAend}}}$ at the time t_{RGAend} when the RGA terminates will then form the queue X^q.

The proof that Algorithm 1 a.a.s. produces a good partial embedding with the properties required in Lemma 3.1 is now quite straightforward: most of the work is to check that the conditions of the various lemmas on RGAs given in Section 2.4 are met. The critical point is to show that certain invariants (see Claim 3.5 below) are maintained.

Proof of Lemma 3.1. We require

$$\rho < \mu \le \frac{1}{100\kappa\Delta_{R'}}, \quad \varepsilon' \le \frac{\mu\zeta d^\Delta \rho 2^{-4/\rho}}{1000\kappa\Delta^2 4^\Delta \Delta_{R'}}, \quad \varepsilon \le \frac{\varepsilon'}{\kappa d}$$

and

$$p \ge 10 r_1 \varepsilon^{-1} \left(\tfrac{\log n}{n}\right)^{1/\Delta}, \quad n > \kappa r_1^2.$$

First, we apply Lemma 2.28 to obtain an ordering τ of the vertices of H with properties (*a*)–(*c*) of this lemma. With this ordering τ we run the random greedy algorithm, Algorithm 1, and in the following we show that a.a.s. it produces a good partial embedding $\psi_{\text{RGA}} = \psi_{t_{\text{RGAend}}}$ satisfying properties (*RGA* 1)–(*RGA* 5) of Lemma 3.1.

We first show that some invariants are maintained (deterministically) throughout the algorithm. Recall that $\pi_t^*(x) := \pi_t(x) + |J_x|$, where $\pi_t(x)$ denotes the number of the embedded neighbours of x at time t.

Algorithm 1: Random greedy algorithm for random graphs

Input : $G \subseteq \Gamma$ and H with partitions satisfying the General Setup; an ordering τ of X^{main}
$t := 0$;
$\psi_0 := \emptyset$;
$Q_0 := \{x \in V(H) : |I_x| < \frac{1}{2}\mu(d-\varepsilon)^{|J_x|}p^{|J_x|}|V^{\mathrm{main}}(x)|\}$;
repeat
    let $x \in X^{\mathrm{main}} \setminus (\mathrm{Dom}(\psi_t) \cup Q_t)$ be the next vertex in the order τ ;
    choose $v \in A_t^{\mathrm{main}}(x) \setminus B_t(x)$ uniformly at random ;
    $\psi_{t+1} := \psi_t \cup \{x \to v\}$;
    $Q_{t+1} := Q_t$;
    forall $y \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_{t+1})$ **do**
        if $\left(\left|A_{t+1}^{\mathrm{main}}(y)\right| < \frac{1}{2}\mu(d-\varepsilon')^{\pi_{t+1}^*(y)}p^{\pi_{t+1}^*(y)}|V^{\mathrm{main}}(y)|\right)$ **then**
            $Q_{t+1} := Q_{t+1} \cup \{y\}$;
        end
    end
    $t := t+1$;
until $\mathrm{Dom}(\psi_t) \cup Q_t = X^{\mathrm{main}}$;
$t_{\mathrm{RGAend}} := t$

Claim 3.5. *The following hold at each time t in Algorithm 1.*

(INV 1) ψ_t *is a good partial embedding.*

(INV 2) *For each $x \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_t)$, either $x \in Q_t$ or we have*

$$|A_t^{\mathrm{main}}(x)| \geq \tfrac{1}{2}\mu(d-\varepsilon')^{\pi_t^*(x)}p^{\pi_t^*(x)}|V^{\mathrm{main}}(x)|.$$

(INV 3) *When we embed x to create ψ_{t+1}, we do so uniformly at random into a set of size at least $\frac{1}{10}\mu\zeta(dp)^{\pi_t^*(x)}|V(x)|$.*

We remark that the bound given in *(INV 3)* is not sharp, but we give this relaxed bound in order to re-use the same invariants in the proofs of our other two blow-up lemmas.

Proof of Claim 3.5. The invariant *(INV 1)* is maintained by definition of $A_t(x)$ and $B_t(x)$. Observe that at the end of each iteration of the repeat loop of Algorithm 1 the queue is updated by adding precisely any vertices which would fail *(INV 2)*, so that this invariant also holds. Finally, by choice of ε' and by *(INV 2)*, if the vertex x is embedded to create ψ_{t+1}, then we have $|A_t^{\mathrm{main}}(x)| \geq \frac{1}{4}\mu(dp)^{\pi_t^*(x)}|V(x)|$. Now, if all neighbours of x are embedded under ψ_t then by Lemma 2.27(*a*) we have *(INV 3)*, while otherwise we apply Lemma 2.27(*b*) with $D = \Delta$, for which the required assumptions on Γ hold, to bound $|B_t(x)|$. We conclude that *(INV 3)* holds. □

We now begin to verify the conclusions of Lemma 3.1. Property *(RGA 2)*, stating that all vertices in $X_i \cap \mathrm{Dom}(\psi_{\mathrm{RGA}})$ are embedded to V_i^{main}, holds by our RGA.

Claim 3.6. *Property* (*RGA* 1) *holds. Moreover, let x be a buffer vertex and y its neighbour whose embedding creates ψ_{t+1}, then*

$$\left|A_t^{\mathrm{main}}(y)\right| \geq \tfrac{2}{3}(dp - \varepsilon'p)^{\pi_t^*}|V(y)|. \tag{3.1}$$

Proof. We require $\mu \leq 1/(100\kappa\Delta_{R'})$. To show (*RGA* 1), i.e., that all neighbours of buffer vertices are in $\mathrm{Dom}(\psi_{\mathrm{RGA}})$, it is sufficient to prove that none of the neighbours of buffer vertices are in $Q_{t_{\mathrm{RGAend}}}$. However, this is clear. Indeed, suppose that $y_1, \ldots, y_b$ with $b \leq \Delta$ are the neighbours of a buffer vertex x appearing in this order in τ. Suppose the embedding of y_1 creates ψ_{t_1+1}. By (*H* 4), buffer vertices are at distance at least five in H, hence, the neighbours of the vertices have distance at least three, and by (*H* 5) they are not image restricted. Therefore, at time t_1 the available candidate set of each y_j is $A_{t_1}^{\mathrm{main}}(y_j) = V^{\mathrm{main}}(y_j) \setminus \mathrm{Im}(\psi_{t_1})$. The size of $A_{t_1}^{\mathrm{main}}(y_j)$ is by (*H* 3), (*G* 1) and choice of μ at least

$$|V^{\mathrm{main}}(y_j)| - 4\kappa\Delta_{R'}\mu|V(y_j)| = (1 - 3\mu - 4\kappa\Delta_{R'}\mu)|V(y_j)| \geq \tfrac{2}{3}|V(y_j)|.$$

It follows that y_j is not added to Q_t for any $t \leq t_1$. Now the vertices $y_1, \ldots, y_b$ are embedded consecutively by Lemma 2.28(*b*), and since they are neighbours of x, by (*H* 2), into distinct clusters of G. By Definition 2.26 (of bad vertices with respect to ψ and Q), for each time t' with $t_1 \leq t' \leq t$, where y_j is embedded to create ψ_{t+1}, we have

$$|A_{t'}^{\mathrm{main}}(y_j)| \geq \tfrac{2}{3}(dp - \varepsilon'p)^{\pi_{t'}^*(y_j)}|V(y_j)|,$$

which gives (3.1) and that y_j never enters Q_t, as desired. $\qquad\square$

We next verify (*RGA* 3) and (*RGA* 4). To this end observe that because Algorithm 1 preserves (*INV* 3), the conditions of Lemma 2.29 are met. Thus with probability at least $1 - 2^{-n/(\kappa r_1)}$, the following event $\mathcal{E}_{\mathrm{L2.29}}$ holds. For every $i \in [r]$ and $W \subseteq V_i$ with $|W| \geq \rho|V_i|$, the number of vertices $x \in X_i \setminus X_i^*$ such that for some time t (at which x is unembedded) we have $\left|C_t(x) \cap W\right| < (dp - \varepsilon'p)^{\pi_t^*(x)}|W|$ is at most $\rho|X_i|$.

Suppose now that $\mathcal{E}_{\mathrm{L2.29}}$ holds. We first show that then (*RGA* 3) holds. Set $W := V_i^{\mathrm{main}} \setminus \mathrm{Im}(\psi_{t_{\mathrm{RGAend}}})$. By (*H* 3) and (*G* 1), we have $|X_i^{\mathrm{main}}| \leq (1 - 4\mu)|X_i|$ and $|V_i^{\mathrm{main}}| = (1 - 3\mu)|V_i|$. We conclude $|W| \geq \mu|V_i|$. Suppose that $x \in Q_{t_{\mathrm{RGAend}}}$. Then there is a first time t at which $x \in Q_t$. Since we have $A_t^{\mathrm{main}}(x) \supseteq C_t(x) \cap W$, by the construction of Q_t in Algorithm 1, we have

$$|C_t(x) \cap W| < \tfrac{1}{2}\mu(d - \varepsilon')^{\pi_t^*(x)}p^{\pi_t^*(x)}|V^{\mathrm{main}}(x)| < \left((d - \varepsilon')p\right)^{\pi_t^*(x)}|W|$$

so that x satisfies condition (2.6) of Lemma 2.29. Since $|W| > \rho|V_i|$, and because $\mathcal{E}_{\mathrm{L2.29}}$ holds, the number of $x \in X_i^{\mathrm{main}} \cap Q_{t_{\mathrm{RGAend}}}$ which are in $X_i \setminus X_i^*$ is at most $\rho|X_i|$. By (*G* 6), we have $|X_i^*| \leq \rho|X_i|$, so $|X_i^{\mathrm{main}} \cap Q_{t_{\mathrm{RGAend}}}| \leq \rho|X_i| + |X_i^*| \leq 2\rho|X_i|$. By definition of Algorithm 1, the vertices $X_i^{\mathrm{main}} \cap Q_{t_{\mathrm{RGAend}}}$ are precisely the vertices of X_i^{main} not in $\mathrm{Dom}(\psi_{t_{\mathrm{RGAend}}})$, giving (*RGA* 3) as desired.

Property (*RGA* 4) follows from $\mathcal{E}_{\mathrm{L2.29}}$ by observing that $X_i^{\mathrm{buf}} \subseteq X_i$ and that when X_i^{buf} is a degree-b buffer we have $(dp - \varepsilon'p)^{\pi_{t_{\mathrm{RGAend}}}^*(x)} \geq (dp)^b/2$ by choice of ε'.

To complete the proof of Lemma 3.1 we now verify that a.a.s. after Algorithm 1 finishes, also (*RGA* 5) is satisfied. By (*a*) and (*c*) of Lemma 2.28, the set of vertices in $N(X^{\mathrm{buf}})$ which are not neighbours of clique buffer vertices is an initial segment of T vertices of τ. We first show that the embedding of these T

vertices cannot fill up vertex neighbourhoods in G. For this purpose we want to apply Lemma 2.30 to the sequence of good partial embeddings produced by Algorithm 1 on this initial segment. Observe that by Claim 3.6, all vertices in this initial segment satisfy (3.1) and none enters the queue. By (3.1), Lemma 2.27, and the choice of ε', each such y, embedded to create ψ_{t+1}, is embedded into a subset of $C_t(y)$ of size at least

$$\tfrac{2}{3}(dp - \varepsilon'p)^{\pi_t^*}|V(y)| - 20\Delta^2\varepsilon'p^{\pi_t^*(y)}|V(y)| > \tfrac{1}{10}(dp)^{\pi_t^*(y)}|V(y)|.$$

Furthermore, for each such y we have $\pi_t^*(y) \leq \Delta - 2$ by Lemma 2.28(b). Finally, by (H3) we have $\left|\mathrm{Dom}(\psi_T) \cap X_i\right| \leq 4\mu\kappa\Delta_{R'}|X_i|$ for each $i \in [r]$. This justifies that we can apply Lemma 2.30 with $B = \Delta$, and the result is that with probability at least $1 - \exp(-\varepsilon pn/r_1)$, the following event $\mathcal{E}_{\mathrm{L2.30}}$ holds. For each $v \in V_i$ and j such that $ij \in E(R')$, we have

$$\left|N_G(v; V_j^{\mathrm{main}}) \setminus \mathrm{Im}(\psi_T)\right| \geq \tfrac{1}{2}\deg_G(v; V_j). \tag{3.2}$$

We assume from now on that X_i^{buf} is a non-clique degree-b buffer, and we fix a vertex $v \in V_i$. We would like to estimate the probability that (RGA5) fails for v. To that end, first fix $x \in X_i^{\mathrm{buf}}$, and let $y_1, \ldots, y_b$ be an enumeration of $N_H(x)$ in the order τ. We now justify that if $\mathcal{E}_{\mathrm{L2.30}}$ holds, then we are in a position to apply Lemma 2.31 with $B = \Delta$. Recall that no vertices of $N(X^{\mathrm{buf}})$ enter the queue, by Claim 3.6, so that $\psi_{\tau(y_1)}$ is the good partial embedding created by the embedding of y_1 (and so on). By (H4), and since $\mathrm{Dom}(\psi_{\tau(y_1)-1}) \subseteq N(X^{\mathrm{buf}})$, no vertices at distance two or less from x in H are embedded in $\psi_{\tau(y_1)-1}$. Note that Lemma 2.28(b) states that if $b = \Delta$ then $y_{b-1}y_b$ is not an edge of H. Algorithm 1 thus creates a sequence $\psi_{\tau(y_1)-1}, \ldots, \psi_{\tau(y_b)}$ of good partial embeddings, and a sequence of queue sets $Q_{\tau(y_1)-1}, \ldots, Q_{\tau(y_b)}$, matching the requirements of Lemma 2.31. Thus, by Lemma 2.31, the probability that $N_H(x)$ is embedded to $N_G(v)$, conditioning on $\psi_{\tau(y_1)-1}$ and that $\left|N_G(v; V_j^{\mathrm{main}}) \setminus \mathrm{Im}(\psi_{\tau(y_1)-1})\right| \geq \tfrac{1}{2}\deg_G(v; V_j)$ for each j such that $ij \in R'$ (which is justified by (3.2)), is at least $(d^\Delta p/10)^b$.

Now we use this to estimate the cumulative effect of all $x \in X_i^{\mathrm{buf}}$. Let $x_1, \ldots, x_m$ be an enumeration of X_i^{buf} according to the order on $N(X^{\mathrm{buf}})$ given by τ. Let $Y_1, \ldots, Y_m$ be Bernoulli random variables with $Y_j = 1$ if either $N_H(x_j)$ is embedded by $\psi_{t_{\mathrm{RGAend}}}$ to $N_G(v)$ or we witness a failure of $\mathcal{E}_{\mathrm{L2.30}}$ before the first neighbour of x_j is embedded. Then we have just shown that Y_j has probability at least $(d^\Delta p/10)^b$ of being one, conditioned on the history up to, but not including, the embedding of the first vertex of $N_H(x_j)$. This history determines Y_{j-1}, so we can apply Lemma 2.2, with $x = 4\mu|X_i|(d^\Delta p/10)^b$ and $\delta = \tfrac{1}{2}$, to conclude that $Y_1 + \cdots + Y_m \geq 2\mu|X_i|(d^\Delta p/10)^b$ with probability at least $1 - \exp(-\mu|X_i|(d^\Delta p/10)^b/3)$. Taking the union bound over all $v \in V(G)$ we see that with probability at least $1 - n\exp(-\varepsilon p^\Delta n/r_1) > 1 - 1/n$ (where the inequality is by choice of p and since $b \leq \Delta$), either we witness a failure of $\mathcal{E}_{\mathrm{L2.30}}$, or ($RGA$5) holds.

Putting together the three probability bounds, we conclude that a.a.s. the events $\mathcal{E}_{\mathrm{L2.29}}$, $\mathcal{E}_{\mathrm{L2.30}}$ and (RGA5) all hold, proving Lemma 3.1. $\qquad\square$

3.3 Queue embedding

In this section we prove Lemma 3.2. The idea is as follows. For each $i \in [r]$ in succession, we embed X_i^{q} into V_i^{q}, maintaining a good partial embedding. The way we do this is as follows. We need to embed

each $x \in X_i^q$ into $C^q(x) \setminus B(x)$, where $B(x)$ is the set of bad vertices for x with respect to the current good partial embedding and $Q = V(H)$. (This choice of Q is made only so that we are able to apply Lemma 2.27 without having to verify the at this time pointless badness condition (2.5).) We therefore try to find a system of distinct representatives $v_x \in C^q(x) \setminus B(x)$ for each $x \in X_i^q$, which we do by verifying Hall's condition. In turn, we verify Hall's condition by showing that its failure implies the existence of a dense spot in G, violating $\mathrm{CON}(\rho, r_1, \Delta)$, which Γ satisfies. At this point one might be concerned that embedding $x' \in X_i^q$ to $v_{x'}$ could change $B(x)$. Observe, however, that since x and x' are at distance at least ten in H by $(H\,2)$ this does not occur.

Proof of Lemma 3.2. We require

$$\mu \leq \frac{1}{6}, \qquad \rho \leq \frac{\mu \zeta d^\Delta}{200\kappa\Delta}, \qquad \varepsilon' \leq \frac{\mu \zeta d^\Delta}{1000\kappa 4^\Delta \Delta^2} \qquad \text{and} \qquad \varepsilon \leq \frac{\varepsilon'}{\kappa d}.$$

Let ψ_0 be a good partial embedding whose image is disjoint from V^q, and suppose that for each $i \in [r]$ the set X_i^q has size at most $2\rho|X_i|$. We now define a sequence of good partial embeddings $\psi_1, \dots, \psi_r$ with $\mathrm{Dom}(\psi_r) = \mathrm{Dom}(\psi_0) \cup \bigcup_{t=1}^r X_t^q$ and $\psi_t(x) \in V_t^q$ for each $x \in X_t^q$ and $1 \leq t \leq r$. We let $C_t(x)$ be the candidate set of x with respect to ψ_t, and so on, and let $B_t(x)$ be the set of bad vertices for x with respect to ψ_t and $Q = V(H)$.

Suppose that for some $1 \leq t \leq r$ we have constructed ψ_{t-1} as above. We let $(v_x)_{x \in X_t^q}$ be a system of distinct representatives for the sets $\left(C_{t-1}^q(x) \setminus B_{t-1}(x) \right)_{x \in X_t^q}$, and we set $\psi_t = \psi_{t-1} \cup \{x \to v_x \colon x \in X_t^q\}$. We need to prove that this system of distinct representatives exists, and that the resulting ψ_t is a good partial embedding.

To see that the system of distinct representatives exists, we verify Hall's condition. Let $X \subseteq X_i^q$ be non-empty, and let $W = \bigcup_{x \in X} C_{t-1}^q(x) \setminus B_{t-1}(x)$. Then we need to show $|W| \geq |X|$. Assume for a contradiction that $|W| < |X|$ holds. By averaging, there is $b \in \{0, \dots, \Delta\}$ such that we find a subset X_b of (not necessarily all) vertices x in X with $\pi_{t-1}^*(x) = b$ of size exactly $\frac{1}{\Delta+1}|W|$. Now each $x \in X_b$ has $C_{t-1}^q(x) \geq (1 - \varepsilon')\mu(dp - \varepsilon'p)^{b-|J_x|}|I_x|$ by $(GPE\,3)$, and $|I_x| \geq \zeta(dp - \varepsilon'p)^{|J_x|}|V_i|$ by $(G\,5)$. We conclude, by Lemma 2.27, that

$$\left| C_{t-1}(x) \setminus B_{t-1}(x) \right| \geq \tfrac{1}{2}\mu\zeta(dp)^b|V_i| - 20\Delta^2\varepsilon'p^b|V_i| \geq \tfrac{1}{4}\mu\zeta(dp)^b|V_i|$$

by choice of ε'. In particular, we see that for each $x \in X_b$ we have $\left| U_{t-1}(x) \cap W \right| \geq \tfrac{1}{4}\mu\zeta(dp)^b|V_i|$. We therefore have

$$\sum_{x \in X_b} \left| U_{t-1}(x) \cap W \right| \geq |X_b| \cdot \tfrac{1}{4}\mu\zeta(dp)^b|V_i|.$$

Since we would like to resort to the congestion property $\mathrm{CON}(\rho, r_1, \Delta)$ we eventually need to pass to a subset W' of W of size $|X_b|$. Picking uniformly at random a subset W' of W of size $|X_b| = \frac{|W|}{\Delta+1}$, we have

$$\mathbb{E}_{W'} \sum_{x \in X_b} \left| U_{t-1}(x) \cap W' \right| \geq \tfrac{|X_b|}{\Delta+1} \cdot \tfrac{1}{4}\mu\zeta(dp)^b|V_i|,$$

and so in particular there is a subset W' of W of size $|X_b|$ such that

$$\sum_{x \in X_b} \left| U_{t-1}(x) \cap W' \right| \geq \tfrac{|X_b|}{\Delta+1} \cdot \tfrac{1}{4}\mu\zeta(dp)^b|V_i|. \tag{3.3}$$

Recall that $\Pi_{t-1}(x)$ is the set of embedded neighbours of x in ψ_{t-1}. We now apply $\text{CON}(\rho, r_1, \Delta)$ with $\mathcal{F} = \left\{ \psi_{t-1}(\Pi_{t-1}(x) \cup J_x) : x \in X_b \right\}$ and the set W'. Note that $|W'| = |X_b| = |\mathcal{F}| \le 2\rho|X_i|/(\Delta+1) \le \rho|V(\Gamma)|$ because the sets $\psi_{t-1}(\Pi_{t-1}(x) \cup J_x)$ are disjoint, so that we can do this. We conclude that the congestion graph $\text{CG}(\Gamma, W', \mathcal{F})$ satisfies

$$e\big(\text{CG}(\Gamma, W', \mathcal{F})\big) \le 7p^b|W'||\mathcal{F}| + \rho p^b n|\mathcal{F}|/r_1. \tag{3.4}$$

But the edges of this congestion graph are precisely the pairs $F_x u$ such that $F_x = \psi_{t-1}(\Pi_{t-1}(x) \cup J_x)$ for $x \in X_b$ and $u \in U_{t-1}(x) \cap W'$, so combining (3.3) and (3.4) we have

$$\tfrac{1}{4(\Delta+1)} \mu \zeta (dp)^b |X_b||V_i| \le 14\rho p^b |X_b||V_i| + \rho \kappa p^b |X_b||V_i|,$$

where we use $|X_b| = |W'| = |\mathcal{F}| \le 2\rho|V_i|$ and $|V_i| \ge n/(\kappa r_1)$. This is a contradiction by choice of ρ. We conclude that the desired system of distinct representatives exists.

Now we show that ψ_t is a good partial embedding. Since the v_x are distinct, ψ_t is injective. Since $C_{t-1}^q(x) \subseteq I_x$ for each x, we have *(GPE 1)* for each $x \in X_i^q$. If *(GPE 3)* or *(GPE 2)* were to fail for some y, then $y \in N_H(x)$ for some $x \in X_i^q$. Since by *(H 2)*, vertices of X_i^q are at distance at least 10 in H, this x is unique. But by definition of $B_{t-1}(x)$ the vertex v_x is not bad for x with respect to ψ_{t-1}, i.e. this case does not occur. Finally, if *(GPE 4)* fails for some $yz \in E(H)$, then again at least one of y and z (possibly both) is a neighbour of some $x \in X_i^q$. Again by *(H 2)* this x is unique, and again by definition of $B_{t-1}(x)$ the vertex v_x is not bad for x with respect to ψ_{t-1}, so this case too does not occur. Thus ψ_t is a good partial embedding as desired.

By induction on t, the final $\psi_q := \psi_r$ is a good partial embedding satisfying the conclusion of Lemma 3.2. □

3.4   Fixing buffer defects

To prove Lemma 3.3 we need to embed the reserved cliques. The basic idea is as follows. If X_i^{buf} is a clique buffer, and a vertex v of V_i is a candidate for too few vertices in X_i^{buf} and not in the image of ψ (we call such a vertex *poor*), then we will embed a reserved clique $K \in \mathcal{K}_i$ on v and some further vertices of V^c. Observe that, since K is a $(\Delta+1)$-clique, none of the vertices of K have embedded neighbours, hence v is in the available candidate set of the vertex of K in X_i.

The only difficulty now is that some of the poor vertices v may lie in V_i^{buf}, and we risk destroying the hard-earned property that every vertex in X_i^{buf} has many candidates. In order to deal with this, we will need to embed some vertices from X_i^{buf} as well, at which point further vertices of V_i may become poor and require embedding, and so on.

Before we prove Lemma 3.3 we state two auxiliary lemmas. The first justifies that the process just described terminates without eating up too many vertices, that is, that we can find for each i small subsets P_i of V_i and D_i of X_i^{buf} such that every vertex $x \in X_i^{\text{buf}} \setminus D_i$ has many candidates in $V_i^{\text{buf}} \setminus P_i$, and every vertex of $V_i \setminus P_i$ is a candidate for many vertices of $X_i^{\text{buf}} \setminus D_i$.

Lemma 3.7. *We assume the General Setup. Suppose that Γ has $\text{CON}(\rho, r_1, \Delta)$ and let i be fixed. Suppose ψ is a good partial embedding whose domain contains $N(X^{\text{buf}})$, and that X_i^{buf} is a clique buffer. Suppose furthermore that all but at most $\rho|V_i|$ vertices of V_i are candidates for at least $\mu(dp)^\Delta|X_i|/2$ vertices of X_i^{buf}. Then there are subsets P_i of V_i and D_i of X_i^{buf} with the following properties.*

(PD 1) *We have $|P_i| < 2\rho|V_i|$ and $|D_i| < 2\rho|X_i|$.*

(PD 2) *Each vertex of $X_i^{\mathrm{buf}} \setminus D_i$ has at least $\mu(dp)^\Delta|V_i|/4$ candidates in $V_i^{\mathrm{buf}} \setminus P_i$.*

(PD 3) *Each vertex of $V_i \setminus P_i$ is a candidate for at least $\mu(dp/100)^\Delta|X_i|$ of the vertices in $X_i^{\mathrm{buf}} \setminus D_i$.*

(PD 4) *Each vertex in D_i has at least $\mu(dp)^\Delta|V_i|/4$ candidates in $P_i \cap V_i^{\mathrm{buf}}$.*

The conclusion (PD 4) will be used later to show that we can embed D_i into $P_i \cap V_i^{\mathrm{buf}}$.

Proof of Lemma 3.7. We require

$$\rho < \frac{\mu d^\Delta}{250\kappa}.$$

We start with $D_i = \emptyset$ and P_i being the set of vertices which fail (PD 3). We sequentially add vertices to D_i and P_i which fail (PD 2) and (PD 3) respectively (note that the property of failing either condition is monotone), until either there are no failing vertices to add to either set or one of the two sets reaches $2\rho|V_i|$ vertices. Observe that any vertex x of D_i has by (GPE 3) at least $(1-\varepsilon')\mu(d-\varepsilon')^\Delta p^\Delta|V_i|$ candidates in V_i^{buf}. Since x must have violated (PD 2) at some point, by choice of ε' it therefore has at least $\mu(dp)^\Delta|V_i|/4$ candidates in $P_i \cap V_i^{\mathrm{buf}}$, establishing (PD 4).

In the process of building D_i and P_i we call a vertex $v \in V_i \setminus P_i$ *poor* if v is a candidate for less than $\mu(dp/100)^\Delta|X_i|$ vertices of $X_i^{\mathrm{buf}} \setminus D_i$. If P_i reaches $2\rho|V_i|$ vertices first, let $P = P_i$ and let D a superset of D_i, both of size exactly $2\rho|V_i|$. Since, by the assumption of the lemma, at most $\rho|V_i|$ vertices in V_i are poor with respect to X_i^{buf}, at least $\rho|V_i|$ of the vertices in P were added to P_i because they became poor, that is, they are candidate for at least $\mu(dp)^\Delta|X_i|/2$ vertices of X_i^{buf}, but for less than $\mu(dp/100)^\Delta|X_i|$ of the vertices $X_i^{\mathrm{buf}} \setminus D_i$. It follows that each is candidate for at least $\mu(dp)^\Delta|X_i|/4$ vertices of D_i, and so there are at least $\rho|V_i|\mu(dp)^\Delta|X_i|/4$ edges in the candidate graph between P and D. Now

$$\begin{aligned}
\tfrac{1}{4}\rho\mu(dp)^\Delta|X_i||V_i| &= \tfrac{1}{8}\rho\mu(dp)^\Delta|X_i||V_i| + \tfrac{1}{8}\rho\mu(dp)^\Delta|X_i||V_i| \\
&> \frac{\mu d^\Delta}{32\rho}p^\Delta|P||D| + \frac{\mu d^\Delta}{16\kappa r_1}p^\Delta n|D| \\
&> 7p^\Delta|P||D| + \rho p^\Delta n|D|/r_1,
\end{aligned}$$

where the first inequality comes from the sizes of P and D, and the second from the choice of ρ sufficiently small. The last line is in contradiction to the congestion condition $\mathrm{CON}(\rho, r_1, \Delta)$, since the candidate graph between P and D is a subgraph of $\mathrm{CG}(\Gamma, P, \mathcal{F}_D)$ with $\mathcal{F}_D := \{\psi(N_H(x)) : x \in D\}$.

If on the other hand D_i reaches $2\rho|V_i|$ vertices first, we define similarly P a superset of P_i and $D = D_i$, both of size $2\rho|V_i|$. Each vertex of D has at least $\mu(dp)^\Delta|V_i|/4$ candidates in P by (PD 4), and thus the candidate graph between P and D has at least $2\rho|X_i|\mu(dp)^\Delta|V_i|/4$ edges, which is larger than in the previous case and so also gives a contradiction to $\mathrm{CON}(\rho, r_1, \Delta)$.

We conclude that the process terminates for lack of failing vertices, which establishes (PD 1). $\qquad\square$

The second lemma simply gives a condition on a vertex $v \in V_i$ under which we can find a clique $K_{\Delta+1}$ containing v in G whose remaining vertices lie in $V^c \setminus \mathrm{Im}(\psi)$.

Lemma 3.8. *We assume the General Setup. Suppose that Γ has $\mathrm{NS}(\varepsilon, r_1, \Delta)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$ and $\mathrm{CON}(\rho, r_1, \Delta)$. Let $i, j_1, \ldots, j_\Delta$ form a copy of $K_{\Delta+1}$ in R' and v be a vertex of V_i. Suppose that ψ is any good partial embedding which embeds no vertex on v, and is such that v has at least*

$$\tfrac{\mu d}{8} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right) \tag{3.5}$$

neighbours in $V_{j_k}^{\mathrm{c}} \setminus \mathrm{Im}(\psi)$ in G for each $1 \le k \le \Delta$. Suppose that there is a copy of $K_{\Delta+1}$ in H using a vertex of $\tilde{X}_i$ and a vertex of each of $X_{j_1}, \ldots, X_{j_\Delta}$. Then there are vertices $v_k \in V_{j_k}^{\mathrm{c}} \setminus \mathrm{Im}(\psi)$ for $1 \le k \le \Delta$, such that $v, v_1, \ldots, v_\Delta$ form a copy of $K_{\Delta+1}$ in G.

Proof. We require

$$\varepsilon < \varepsilon' < \frac{(d/32)^\Delta \mu}{2\kappa\Delta^2}.$$

We choose the vertices $v_1, \ldots, v_\Delta$ in succession. When we choose v_s, if $s < \Delta$ we require that for each $s+1 \le k \le \Delta$ we have

$$|N_\Gamma^*(v, v_1, \ldots, v_s) \cap V_{j_k}| = (p \pm \varepsilon p)^s \deg_\Gamma(v; V_{j_k}) \text{ and} \tag{3.6}$$

$$|N_G^*(v, v_1, \ldots, v_s) \cap \left(V_{j_k}^{\mathrm{c}} \setminus \mathrm{Im}(\psi)\right)| \ge \left(\tfrac{dp}{4}\right)^s \tfrac{\mu d}{8} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right). \tag{3.7}$$

If $s < \Delta - 1$ we further require that for each $s+1 \le k < k' \le \Delta$ the pair

$$\left(N_\Gamma^*(v, v_1, \ldots, v_s) \cap V_{j_k}, N_\Gamma^*(v, v_1, \ldots, v_s) \cap V_{j_{k'}}\right) \tag{3.8}$$

is $(\varepsilon_{s,s}, d, p)$-regular.

Note that these conditions are satisfied, with $s = 0$, before we have chosen any vertices. Indeed, the first is a tautology, the second is the assumption on v in the lemma statement, and the third is a statement that V_i has two-sided inheritance with respect to V_{j_k} and $V_{j_{k'}}$, which holds since we assumed there is a triangle of X using one vertex of each of $\tilde{X}_i$, X_{j_k} and $X_{j_{k'}}$ and by *(H 1)* and *(G 2)*.

Suppose we have chosen vertices $v_1, \ldots, v_{s-1}$ so far and that we have $s \le \Delta - 1$. Because Γ has $\mathrm{NS}(\varepsilon, r_1, \Delta)$, at step s the number of vertices w failing condition (3.6) with $v_s = w$, is at most

$$\Delta \varepsilon p^{\Delta-1} |V(\Gamma)|/r_1^2 < \left(\tfrac{dp}{4}\right)^{s-1} \tfrac{\mu d}{32} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right).$$

By $(\varepsilon_{s-1,s-1}, d, p)$-regularity of the pair

$$\left(N_\Gamma^*(v, v_1, \ldots, v_{s-1}) \cap V_{j_k}, N_\Gamma^*(v, v_1, \ldots, v_{s-1}) \cap V_{j_{k'}}\right),$$

and using the upper bound from (3.6) on $|N_\Gamma^*(v, v_1, \ldots, v_{s-1}) \cap V_{j_k}|$, at most

$$\Delta \varepsilon' (p + \varepsilon p)^{s-1} \deg_\Gamma(v; V_{j_k}) < \left(\tfrac{dp}{4}\right)^{s-1} \tfrac{\mu d}{32} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right)$$

vertices $w \in V_{j_s}^{\mathrm{c}} \setminus \mathrm{Im}(\psi)$ fail condition (3.7) with $v_s = w$. Finally, if $s \le \Delta - 2$ and since Γ has $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$, at most

$$\Delta^2 \varepsilon p^{\Delta-2} |V(\Gamma)|/r_1^2 < \left(\tfrac{dp}{4}\right)^{s-1} \tfrac{\mu d}{32} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right)$$

vertices fail condition (3.8). Thus there are at least

$$\left(\tfrac{dp}{4}\right)^{s-1} \tfrac{\mu d}{32} \max\left(p|V_{j_k}|, \deg_\Gamma(v; V_{j_k})\right) > 0$$

vertices from $N_G^*(v, v_1, \ldots, v_{s-1}) \cap \left(V_{j_s}^c \setminus \mathrm{Im}(\psi)\right)$ which satisfy all three conditions, and we can choose v_s for each $1 \leq s \leq \Delta$ as desired.

Furthermore observe that for $s = \Delta$ it suffices to choose an arbitrary vertex from $N_G^*(v, v_1, \ldots, v_{\Delta-1}) \cap \left(V_{j_\Delta}^c \setminus \mathrm{Im}(\psi)\right)$. The lemma follows. $\qquad\square$

We can now prove Lemma 3.3. We construct the desired embedding in two steps. First, we embed each D_i into $P_i \cap V_i^{\mathrm{buf}}$, which by verifying Hall's condition. Let P_i' be those vertices of P_i not covered in this process. It then remains to cover all P_i' using the reserved cliques, and embed any left-over reserved cliques at the end of this process. We do this sequentially by applying Lemma 3.8 to find a destination for a reserved clique, either letting v be the next poor vertex or, if all of these are used, letting v be some vertex of V^c. Note that in this step, we need to maintain the degree condition (3.5) of Lemma 3.8. We have this condition initially by $(G\,3)$, and we will see that to maintain it, it suffices to choose at each step a most dangerous vertex v, that is, one minimising the parameter mindeg we define in the proof.

Proof of Lemma 3.3. We require

$$\rho < \min\left(\frac{\mu d^\Delta}{250\kappa}, \frac{\mu d}{64\kappa(\Delta_{R'} + 1)}\right) \quad \text{and } \varepsilon < \varepsilon' < \frac{d^\Delta}{32^\Delta \kappa \Delta \Delta_{R'}},$$

and assume that Γ has $\mathrm{NS}(\varepsilon, r_1, \Delta)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$ and $\mathrm{CON}(\rho, r_1, \Delta)$.

For each i such that X_i^{buf} is a clique buffer, we let P_i and D_i be as given by Lemma 3.7. If X_i^{buf} is not a clique buffer, then we let $P_i = D_i = \emptyset$.

Now for each i in succession, we find a matching in the candidate graph between D_i and $P_i \cap V_i^{\mathrm{buf}}$. To do this we simply verify Hall's condition, using the congestion condition and *(PD 4)*. Specifically, suppose we have a non-empty $W \subseteq D_i$, and $Z \subseteq P_i$ is the set of vertices in P_i which are candidate for some $w \in W$. We need to verify that $|W| \leq |Z|$. If this were false, then we could let Z' be a superset of Z of size $|W|$, and the number of edges in the candidate graph between W and Z' is by *(PD 4)* at least $\mu(dp)^\Delta |V_i||W|/4$. Notice that $|D_i|, |P_i| < 2\rho|V_i|$ by *(PD 1)*. Hence we have

$$\begin{aligned}
\tfrac{1}{4}\mu(dp)^\Delta |V_i||W| &= \tfrac{1}{8}\mu(dp)^\Delta |V_i||W| + \tfrac{1}{8}\mu(dp)^\Delta |V_i||W| \\
&> \frac{\mu d^\Delta}{16\rho} p^\Delta |Z'||W| + \frac{\mu d^\Delta}{8\kappa r_1} p^\Delta n|W| \\
&> 7 p^\Delta |Z'||W| + \rho p^\Delta n|W|/r_1,
\end{aligned}$$

which is in contradiction to $\mathrm{CON}(\rho, r_1, \Delta)$, applied with Z' and the family $\mathcal{F} = \{\psi(N_H(w)) \colon w \in W\}$ to the congestion graph $\mathrm{CG}(\Gamma, Z', \mathcal{F})$. We conclude that Hall's condition holds and hence the desired matching exists. We let ψ' be the embedding obtained from ψ by embedding each vertex of each D_i to its matching partner in $P_i \cap V_i^{\mathrm{buf}}$, and let $P_i' \subseteq P_i$ be the set of those vertices in P_i not embedded by ψ'. Observe that ψ' is a good partial embedding since all neighbours of the buffer vertices are already embedded under ψ, and thus *(GPE 1)*–*(GPE 4)* trivially hold for ψ'.

Algorithm 2: Removing poor vertices

$t := 0;$
$\psi_0 := \psi';$
while $\bigcup_i \mathcal{K}_i \neq \emptyset$ **do**
$\quad \big| \quad I := \{i : \text{there is an unembedded clique in } \mathcal{K}_i\};$
$\quad \big| \quad$ **if** $\bigcup_i P_i' \setminus \text{Im}(\psi_t) \neq \emptyset$ **then**
$\quad \big| \quad \big| \quad$ choose $v \in \bigcup_i P_i' \setminus \text{Im}(\psi_t)$ minimising $\text{mindeg}_{\psi_t}(v)$;
$\quad \big| \quad$ **else**
$\quad \big| \quad \big| \quad$ choose $v \in \bigcup_{i \in I} V_i^c \setminus \text{Im}(\psi_t)$ minimising $\text{mindeg}_{\psi_t}(v)$;
$\quad \big| \quad$ **end**
$\quad \big| \quad$ set i such that $v \in V_i$;
$\quad \big| \quad$ choose an unembedded clique K_H in $\mathcal{K}_i$;
$\quad \big| \quad$ choose a clique K_G containing v and one vertex from each $V_j^c \setminus \text{Im}(\psi_t)$ such that X_j intersects
$\quad \big| \qquad K_H$ and $j \neq i$;
$\quad \big| \quad$ set $\psi_{t+1} := \psi_t \cup \{K_H \to K_G\};$
$\quad \big| \quad t := t + 1;$
end
return $\psi_{\text{good}} := \psi_t$;

As second step, we provide an algorithm, Algorithm 2, to embed the reserved cliques in $\mathcal{K}_i$ for each $i \in [r_1]$. Given a vertex $v \in V_i$ and a partial embedding ψ we define the parameter

$$\text{mindeg}_\psi(v) := \min_{j:ij \in R'} \frac{\deg_G\left(v; V_j^c \setminus \text{Im}(\psi)\right)}{\max\left(p|V_j|, \deg_\Gamma(v; V_j)\right)},$$

which is used in Algorithm 2. Observe that we have initially $\text{mindeg}_\psi(v) \geq \mu(d - \varepsilon)/2$ for all $v \in V_i$ with $\deg_{R'}(i) > 0$, which follows from $(G\,3)$. We claim that the algorithm runs correctly and the finally returned ψ_{good} is the desired good partial embedding. It remains to justify this claim.

Roughly speaking each of the vertices from P_i' can serve us as an image for each of the $2\rho|X_i|$ cliques from $\mathcal{K}_i$. During the first phase of the algorithm, that is, while the if-condition is true, we embed reserved cliques in these vertices, always choosing a vertex v minimising $\text{mindeg}_\psi(v)$, which will guarantee successful completion. Note that $|P_i'| < 2\rho|V_i|$ holds by $(PD\,1)$ and $|\mathcal{K}_i| = 2\rho|X_i|$, so if the algorithm doesn't fail then all the vertices from $\bigcup_i P_i'$ will be used. In the second phase, when the if-condition becomes false, we have to take care of the remaining unembedded cliques in $\bigcup_i \mathcal{K}_i$. Since the algorithm embeds only onto vertices $\bigcup P_i' \cup V_i^c$, if it completes successfully, then the desired properties $(BD\,2)$ and $(BD\,3)$ are guaranteed for ψ_{good} by $(PD\,2)$ and $(PD\,3)$.

Therefore, it remains to show that the algorithm succeeds, that is, that there is no failure while choosing a clique K_G in some iteration. By Lemma 3.8, this is guaranteed if we can show that at each time t the vertex $v \in V_i$ chosen has at least

$$\frac{\mu d}{8} \max\left(p|V_j|, \deg_\Gamma(v; V_j)\right) \tag{3.9}$$

neighbours in each set $V_j^c \setminus \text{Im}(\psi_t)$ such that $ij \in R'$.

By the assumption of the lemma, for each i, the set X_i^c of vertices $x \in X_i$ contained in some reserved clique from $\bigcup_j \mathcal{K}_j$ satisfies $|X_i^c| \le 2\kappa(\Delta_{R'}+1)\rho|X_i|$. It follows that for any t and j we have

$$|\text{Im}(\psi_t) \cap V_j^c| \le 2\kappa(\Delta_{R'}+1)\rho|X_j|. \tag{3.10}$$

Claim 3.9. *For any t at most $\Delta_{R'}\varepsilon p^{\Delta-1}|V(\Gamma)|/r_1^2$ vertices $v \in V_i$ satisfy*

$$\text{mindeg}_{\psi_t}(v) < 3\mu d/16.$$

Proof. Each vertex v of V_i has at least

$$\mu(d-\varepsilon)\max\big(p|V_j|, \deg_\Gamma(v;V_j)/2\big) > \tfrac{\mu d}{4}\max\big(p|V_j|, \deg_\Gamma(v;V_j)\big)$$

neighbours in V_j^c by (G 3). Moreover, it is a direct consequence of the neighbourhood size property $\text{NS}(\varepsilon, r_1, \Delta)$ of Γ and (3.10) that the number of vertices in Γ which have more than $4\kappa(\Delta_{R'}+1)\rho p|X_j|$ neighbours in $\text{Im}(\psi_t) \cap V_j^c$ is at most $\varepsilon p^{\Delta-1}|V(\Gamma)|/r_1^2$. Since $4\kappa(\Delta_{R'}+1)\rho p|X_j| < \mu d p|V_j|/16$, it follows that for any t at most $\Delta_{R'}\varepsilon p^{\Delta-1}|V(\Gamma)|/r_1^2$ vertices in V_i have fewer than $\tfrac{3\mu d}{16}\max\big(p|V_j|, \deg_\Gamma(v;V_j)\big)$ neighbours in any set $V_j^c \setminus \text{Im}(\psi_t)$ such that $ij \in R'$, that is, they have $\text{mindeg}_{\psi_t}(v) < 3\mu d/16$. □

In the remainder we show that if Algorithm 2 chooses a vertex v such that (3.9) fails, then we obtain a contradiction to Claim 3.9. Indeed, in each iteration at most one vertex is embedded in any given V_j^c. So if there is a time t at which a vertex $v \in V_i$ is chosen with fewer neighbours than the bound in (3.9) in some set $V_j^c \setminus \text{Im}(\psi_t)$ such that $ij \in R'$, then at each of the preceding $\mu dp|V_j|/16$ times t' the vertex v had $\text{mindeg}_{\psi_{t'}}(v) < 3\mu d/16$. Hence each of the at least $\mu dpn/(16\kappa r_1)$ vertices v' chosen at these times has

$$\text{mindeg}_{\psi_t}(v') \le \text{mindeg}_{\psi_{t'}}(v') < 3\mu d/16.$$

Since there are at most r_1 clusters V_i in G, and since

$$\mu dpn/(16\kappa r_1) > r_1\Delta_{R'}\varepsilon p^{\Delta-1}|V(\Gamma)|/r_1^2,$$

this contradicts Claim 3.9. □

Chapter 4

Proof of the blow-up lemma for bijumbled graphs

4.1 The RGA lemma and the proof of the blow-up lemma

In this section we prove Lemma 1.25, conditional on the random greedy lemma (Lemma 4.1) which we prove in Section 4.2.

The proof is quite similar to that of Lemma 1.21, and indeed we make use of the same General Setup. The difference is that we can no longer use property $\mathrm{CON}(\rho, r_1, \Delta)$ as a bijumbled graph Γ need not satisfy it (see the discussion in Section 2.2). Instead, we 'replace' it with the lopsided neighbourhood size property $\mathrm{LNS}(\varepsilon, r_1, \Delta)$.

In the next section we will use a somewhat different random greedy strategy than Algorithm 1 to prove the following lemma. However, we will re-use the auxiliary lemmas proved in Section 2.4 in its proof (recall that these lemmas are not dependent on Algorithm 1).

Lemma 4.1 (bijumbled RGA lemma). *We assume the General Setup. Suppose further that Γ has properties $\mathrm{NS}(\varepsilon, r_1, \Delta + 1)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta + 1)$ and $\mathrm{LNS}(\varepsilon, r_1, \Delta)$, and that at most $\rho p^\Delta |X_i|$ vertices in each X_i are image restricted. Then there is a good partial embedding ψ_{RGA} such that the following hold for each i. Let b be such that X_i^{buf} is a degree-b buffer.*

(PRGA 1) *Every vertex in X^{main} is embedded to $V^{\mathrm{main}} \cup V^q$ by ψ_{RGA}, and no vertex in X^{buf} is embedded.*

(PRGA 2) *Every vertex in V_i is a candidate for at least $\mu (d^\Delta p / 100)^b |X_i|$ vertices of X_i^{buf}.*

(PRGA 3) *For every set $W \subseteq V_i$ of size at least $\rho |V_i|$, there are at most $\rho |X_i|$ vertices in X_i^{buf} with fewer than $\frac{1}{2}(dp)^b |W|$ candidates in W.*

(PRGA 4) *For all $x \in X_i^{\mathrm{buf}}$, letting*

$$Y := \left\{ y \in X_i^{\mathrm{buf}} : \left| U^{\mathrm{buf}}(y) \cap C^{\mathrm{buf}}(x) \right| > (p + \varepsilon p)^b |C^{\mathrm{buf}}(x)| \right\},$$

we have

$$\sum_{y \in Y} \frac{\left| U^{\mathrm{buf}}(y) \cap C^{\mathrm{buf}}(x) \right|}{\left| C^{\mathrm{buf}}(y) \right|} \leq \frac{|C^{\mathrm{buf}}(x)|}{20} .$$

86

(*PRGA* 5) *For each i, for all but at most $\varepsilon' p^\Delta |V_i|$ vertices $v \in V_i$, there are at most $\varepsilon' p^\Delta |V_i|$ vertices $v' \in V_i$ such that*

$$\left|\{x \in X_i^{\text{buf}} : v, v' \in C(x)\}\right| > 24\mu\Delta^2 \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^b p^{2b}|X_i|.$$

Comparing this lemma to Lemma 3.1 (RGA lemma), one notices that we demand more of the NS and RI pseudorandomness properties. The reason for doing this is that we can then make use of Lemmas 2.30 and 2.31 for all buffer vertices, whether or not they are in copies of $K_{\Delta+1}$, and thus there is no exception for clique buffers in (*PRGA* 2). This simplifies our proof substantially, as we no longer need to fix buffer defects. We will explain later why this does not harm our eventual bijumbledness requirement.

Property (*PRGA* 3) corresponds to (*RGA* 4). Further, properties (*PRGA* 4) and (*PRGA* 5) form the 'no dense spots' property we require in order to complete the embedding. In Lemma 3.1 no such property appears, but in random graphs the CON property guarantees that no dense spots exist. Somewhat informally, (*PRGA* 4) asserts that typically candidate sets intersect as if they are random sets, while property (*PRGA* 5) states the same about intersections of the neighbourhoods of vertices from $V(G)$ in the candidate graph.

Observe that, in comparison to Lemma 3.1, here we ask for an additional pseudorandomness property $\text{LNS}(\varepsilon, r_1, \Delta)$, which we require to establish (*PRGA* 4) and (*PRGA* 5) and in order to embed all vertices.

4.1.1 Outline of the proof of Lemma 1.25

The proof of Lemma 1.25 now looks quite similar to the proof of Lemma 1.21 (the blow-up lemma for random graphs). Again, we use Lemma 2.22 to obtain the General Setup. However, this time we do not choose any reserved cliques, whether or not we have buffer vertices in cliques. We note that although we could state a version of Lemma 4.1 which asked only for $\text{NS}(\varepsilon, r_1, \Delta)$ and $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta)$ rather than with $\Delta + 1$, and which made exceptions for clique buffers, this would not improve our eventual bijumbledness requirement on Γ, except for $\Delta = 2$, since for $\Delta \geq 3$ the requirement is determined by $\text{LNS}(\varepsilon, r_1, \Delta)$. We apply Lemma 4.1 to obtain a good partial embedding whose domain is X^{main} with the additional properties stated there. We then have only to embed X^{buf}. Again, we do this one part at a time, by verifying Hall's condition, and again we separate the verification into small, large and medium-sized subsets of X_i^{buf}. Again, the 'medium-sized' case is dealt with quickly using (*PRGA* 3).

However, now we approach the 'small' and 'large' cases differently, since we no longer have access to the CON property. In the 'small' case (see Claim 4.2), our approach is the following. Given $X \subseteq X_i^{\text{buf}}$, we construct an auxiliary embedding of X into V_i^{buf} such that each x is embedded into $C(x)$. Although this auxiliary embedding is *not* part of the embedding we will finally give, its existence verifies Hall's condition for X. The way we will construct the embedding is simply to embed X vertex by vertex, at each step choosing for x uniformly at random a so far unused vertex of $C^{\text{buf}}(x)$. Since ψ_{RGA} has (*GPE* 3) we know $C^{\text{buf}}(x)$ is always reasonably large, and we will show that with high probability it is never more than half covered. Let us briefly explain how this works. We assume that no candidate sets are more than half covered before reaching x. When we embedded some vertex y, the probability of embedding it to $C(x)$ was at most

$$\frac{\left|U^{\text{buf}}(y) \cap C^{\text{buf}}(x)\right|}{\frac{1}{2}|C^{\text{buf}}(y)|} \tag{4.1}$$

since (by assumption) we embed y into a set of size at least $\frac{1}{2}|C^{\mathrm{buf}}(y)|$, of which certainly at most $\left|U^{\mathrm{buf}}(y) \cap C^{\mathrm{buf}}(x)\right|$ vertices are in $C^{\mathrm{buf}}(x)$. But *(PRGA 4)* gives us an upper bound for the sum of these probabilities over all $y \in X_i^{\mathrm{buf}}$ such that (4.1) is 'large', and it is easy to account for the $y \in X_i^{\mathrm{buf}}$ such that (4.1) is not large. We can thus apply Lemma 2.2 to conclude that it is unlikely that $C^{\mathrm{buf}}(x)$ was more than half covered, and by the union bound over all $x \in X_i^{\mathrm{buf}}$ we conclude that with high probability we never fail.

For the 'large' case (see Claim 4.3) we use a similar argument, but using *(PRGA 2)* and *(PRGA 5)* instead of *(GPE 3)* and *(PRGA 4)*.

4.1.2 Proof of Lemma 1.25

Now we give the full details of the proof outlined above.

Proof of Lemma 1.25. First we choose constants as follows. Given Δ, $\Delta_{R'}^{\mathrm{BL}}$, Δ_J integers, α^{BL}, ζ^{BL} and $d > 0$, and $\kappa^{\mathrm{BL}} > 1$, we set $\vartheta = \Delta$, $\Delta_{R'} = 8(\Delta + \Delta_J)^{10}\Delta_{R'}^{\mathrm{BL}}$, $\alpha = \frac{1}{2}\alpha^{\mathrm{BL}}$, $\zeta = \frac{1}{2}\zeta^{\mathrm{BL}}$ and $\kappa = 2\kappa^{\mathrm{BL}}$. We now choose

$$\mu = \frac{\alpha d^{\Delta}}{20000\kappa\Delta_{R'}^4\Delta^{10}(\Delta+2)}, \qquad\qquad \rho = \frac{\mu^2\zeta^2}{2000\Delta^2}\left(\frac{\mu\zeta d^{4\Delta+2}}{10^{6\Delta}}\right)^{\Delta},$$

$$\text{and} \qquad\qquad \varepsilon' = \frac{\mu\zeta d^{\Delta^2}\rho}{10^{6\Delta}\Delta^3\kappa 4^{\Delta}}2^{-4/\rho}.$$

Now for input Δ, d and ε', Lemma 2.18 returns constants $\varepsilon_{a,b}$ and $\varepsilon_{\mathrm{L}2.18} > 0$. We set

$$\varepsilon = \min\left(\frac{2^{10\Delta}d(\varepsilon')^2}{\kappa\Delta_{R'}}, \varepsilon_{\mathrm{L}2.18}\right).$$

We let $\varepsilon^{\mathrm{BL}} = \frac{1}{16}(\Delta+\Delta_J)^{-10}\varepsilon$ and $\rho^{\mathrm{BL}} = \frac{1}{16}(\Delta+\Delta_J)^{-10}\rho$.

Now Lemma 1.25 returns $\varepsilon^{\mathrm{BL}}$ and ρ^{BL}. Given r_1^{BL} we let $r_1 = 8(\Delta+\Delta_J)^{10}r_1^{\mathrm{BL}}$. We choose c sufficiently small for Lemma 2.18 with input Δ, d, ε' and $T = r_1$.

We require

$$p \geq \frac{10^6(\Delta+\Delta_J)^{10}\kappa r_1}{\mu\zeta\rho d^{\Delta}\varepsilon^{-2}}\left(\frac{\log n}{n}\right)^{1/(2\Delta)},$$

where the condition on the constant in front of the $\left(\frac{\log n}{n}\right)^{1/(2\Delta)}$-term comes from Lemma 2.22, while the term $\left(\frac{\log n}{n}\right)^{1/(2\Delta)}$ is required for Lemma 4.1. We remark, however, that this lower bound is not a restriction since a somewhat stronger lower bound on p of the form $\Omega(n^{-1/2(t-1)})$ for any $(p, cp^t n)$-bijumbled graph follows by considering the inequality (1.1) which defines bijumbledness for, say, sets $X = \{v\}$ and $Y = V(\Gamma) \setminus (\{v\} \cup N_{\Gamma}(v))$.

Let Γ be an n-vertex, $(p, cp^{\max(4,3\Delta/2+1/2)}n)$-bijumbled graph. Then Lemma 2.18 states that Γ has the three properties $\mathrm{NS}(\varepsilon, r_1, \Delta+1)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta+1)$ and $\mathrm{LNS}(\varepsilon, r_1, \Delta)$. From now on we will simply assume Γ is an n-vertex graph which satisfies these three properties.

Given a graph R_{BL} on $r_{\mathrm{BL}} \leq r_1^{\mathrm{BL}}$ vertices, a spanning subgraph R_{BL}' with $\Delta(R_{\mathrm{BL}}') \leq \Delta_{R'}^{\mathrm{BL}}$, and graphs H and $G \subseteq \Gamma$ with vertex partitions $\mathcal{X}^{\mathrm{BL}}$ and $\mathcal{V}^{\mathrm{BL}}$, a family of potential buffer vertices $\tilde{\mathcal{X}}^{\mathrm{BL}}$, and a $(\rho^{\mathrm{BL}}p^{\Delta}, \zeta^{\mathrm{BL}}, \Delta, \Delta_J)$-restriction pair $\mathcal{J}^{\mathrm{BL}}$, $\mathcal{J}$, suppose that the conditions of Lemma 1.25 are satisfied. Then the good partitions lemma, Lemma 2.22, (with $b_{\mathrm{L}2.22} = \Delta$) gives a graph R on $r \leq r_1$ vertices, a spanning

subgraph R' with $\Delta(R') \leq \Delta_{R'}$, and κ-balanced size-compatible partitions $\mathcal{X}$ and $\mathcal{V}$ of H and G respectively, each part having size at least $n/(\kappa r_1)$, together with a family $\tilde{\mathcal{X}}$ of potential buffer vertices and $\mathcal{I}$ of image restrictions, subsets X_i^{buf} of $\tilde{X}_i$ for each $i \in [r]$, and partitions $V_i = V_i^{\text{main}} \,\dot\cup\, V_i^{\text{q}} \,\dot\cup\, V_i^{\text{c}} \,\dot\cup\, V_i^{\text{buf}}$ for each $i \in [r]$ which satisfy the General Setup.

We let $X^{\text{main}} = V(H) \setminus X^{\text{buf}}$. We now begin the embedding of H into G. By Lemma 4.1, there is a good partial embedding ψ_{RGA} of X^{main} into $V^{\text{main}} \cup V^{\text{q}}$ with properties $(PRGA\,1)$–$(PRGA\,5)$ as stated in that lemma.

Our aim now is to complete the embedding of H into G by finding for each i a matching in the available candidate graph between $V_i' := V_i \setminus \text{Im}(\psi_{\text{RGA}})$ and X_i^{buf}. Let us assume that X_i^{buf} is a degree-b buffer. To show there is a matching, we will verify Hall's condition, so let Y be a non-empty subset of X_i^{buf}. Let U be the set of vertices in V_i' which are candidate for at least one vertex of Y, then our aim is to show $|U| \geq |Y|$. We split this into three cases, the harder two of which are done in the following two claims.

Claim 4.2. *If $0 < |Y| \leq \rho|X_i|$, then we have $|U| \geq |Y|$.*

Proof. By $(H\,5)$ no $y \in Y$ is image restricted, so because ψ_{RGA} is a good partial embedding, by $(GPE\,3)$ for each $y \in Y$ we have

$$|C^{\text{buf}}(y)| \geq (1-\varepsilon')\mu(d-\varepsilon')^b p^b |V_i| \geq \tfrac{1}{2}\mu(dp)^b|V_i|,$$

where the second inequality is by choice of ε'. Let $y_1,\ldots,y_{|Y|}$ be an enumeration of Y. Now for each $j = 1,\ldots,|Y|$ we choose v_j uniformly at random from $C^{\text{buf}}(y_j) \setminus \{v_1,\ldots,v_{j-1}\}$ if this is possible; if not, we say v_j does not exist. Our aim is to show that this does not occur, in other words that we obtain $|Y|$ distinct vertices of $|U|$, verifying the claim.

Let $\mathcal{H}_0$ be the empty history, and for each $1 \leq j \leq |Y|$, let $\mathcal{H}_j$ be the history of this process up to and including the choice of v_j. We claim that a.a.s.

$$\left|C^{\text{buf}}(y_j) \setminus \{v_1,\ldots,v_{j-1}\}\right| \geq \tfrac{1}{2}\left|C^{\text{buf}}(y_j)\right| \tag{4.2}$$

holds for each j. For a given j we define random variables $Z_\ell^{(j)}$, $\ell \in [j-1]$, as follows. We set $Z_\ell^{(j)} = 1$ if $v_\ell \in C^{\text{buf}}(y_j)$ and $|C^{\text{buf}}(y_\ell) \setminus \{v_1,\ldots,v_{\ell-1}\}| \geq \tfrac{1}{2}|C^{\text{buf}}(y_\ell)|$ hold, and $Z_\ell^{(j)} = 0$ otherwise. For any history $\mathcal{H}_{\ell-1}$ of this process we can bound the conditional expectation of $Z_\ell^{(j)}$ by

$$\mathbb{E}(Z_\ell^{(j)}|\mathcal{H}_{\ell-1}) = \mathbb{P}(Z_\ell^{(j)} = 1|\mathcal{H}_{\ell-1})$$
$$\leq \frac{\left|C^{\text{buf}}(y_\ell) \cap C^{\text{buf}}(y_j) \setminus \{v_1,\ldots,v_{\ell-1}\}\right|}{\left|C^{\text{buf}}(y_\ell) \setminus \{v_1,\ldots,v_{\ell-1}\}\right|} \leq \frac{\left|U^{\text{buf}}(y_\ell) \cap C^{\text{buf}}(y_j)\right|}{\tfrac{1}{2}\left|C^{\text{buf}}(y_\ell)\right|}. \tag{4.3}$$

Observe that $\sum_{\ell=1}^{j-1} Z_\ell^{(j)}$ is exactly the number of those vertices v_k among v_1, $\ldots$, $v_{\ell-1}$ which lie in $C^{\text{buf}}(y_j)$ and were chosen from a set of size at least $\tfrac{1}{2}|C^{\text{buf}}(y_\ell)|$. To obtain an upper bound on the expectation of $\sum_{\ell=1}^{j-1} Z_\ell^{(j)}$, we sum up (4.3) for all $\ell \in [j-1]$. We split this sum according to whether

$\left|U^{\text{buf}}(y_\ell)\cap C^{\text{buf}}(y_j)\right| \le (p+\varepsilon p)^b |C^{\text{buf}}(y_j)|$ or not. In the former case, each summand is by *(GPE 2)* and *(GPE 3)* at most

$$\frac{2(p+\varepsilon p)^b \mu(p+\varepsilon p)^b |V_i|}{(1-\varepsilon')\mu(dp-\varepsilon'p)^b |V_i|} \le 4d^{-\Delta} p^b,$$

where the inequality is by choice of ε and ε', while *(PRGA 4)* bounds the sum over the remaining terms. We get

$$\sum_{\ell=1}^{j-1} \mathbb{E}\left[Z_\ell^{(j)}\,\big|\,\mathscr{H}_{\ell-1}\right] \le \sum_{\ell=1}^{j-1} \frac{\left|U^{\text{buf}}(y_\ell)\cap C^{\text{buf}}(y_j)\right|}{\frac{1}{2}\left|C^{\text{buf}}(y_\ell)\right|}$$

$$\le 4(j-1)d^{-\Delta}p^b + \frac{|C^{\text{buf}}(y_j)|}{10} \le \frac{|C^{\text{buf}}(y_j)|}{5}$$

where the last inequality uses $j \le |Y| \le \rho|X_i|$, *(GPE 3)* and the choice of ρ and ε. Using the sequential dependence lemma, Lemma 2.2, with $\delta = 1$, we conclude that the probability of the event $\sum_{\ell=1}^{j-1} Z_\ell^{(j)} > \frac{2}{5}|C^{\text{buf}}(y_j)|$ is at most $e^{-|C^{\text{buf}}(y_j)|/15} \le e^{-\mu d^b p^b |V_i|/30}$.

Observe that the event $\sum_{\ell=1}^{j-1} Z_\ell^{(j)} > \frac{2}{5}|C^{\text{buf}}(y_j)|$ contains the event that (4.2) fails for j, given that it did not fail for smaller j. Thus, taking a union bound over the $|Y| \le n$ values of j, we see that the probability that (4.2) fails at any stage is at most

$$n \cdot e^{-\mu d^b p^b |V_i|/30},$$

which is smaller than one for all sufficiently large n by choice of p.

We conclude that there is a positive probability of choosing $|Y|$ distinct vertices in U, so $|U| \ge |Y|$ as desired. □

As previously mentioned, although the proof of Claim 4.2 constructs an auxiliary embedding, this embedding is not part of the final embedding of H into G, and exists only to verify the claim. The same goes for the following claim.

Claim 4.3. *If $|Y| > |X_i^{\text{buf}}| - \rho|X_i|$, then $|U| \ge |Y|$.*

Proof. The vertices $V_i' \setminus U$ are candidates only for vertices of $X_i^{\text{buf}} \setminus Y$, and since $|V_i'| = |X_i^{\text{buf}}|$, what we need to show is that $|V_i' \setminus U| \le |X_i^{\text{buf}} \setminus Y|$ (as explained in detail in the proof of Lemma 3.4).

Let $v \in V_i'$ be a vertex and let $C(v)$ denote those vertices x from X_i^{buf} for which v is a candidate, i.e. $C(v)$ denotes the neighbours of v in the candidate graph between $V_i' \setminus U$ and $X_i^{\text{buf}} \setminus Y$. By *(PRGA 2)*, for each vertex v of $V_i' \setminus U$ we have $|C(v)| \ge \mu(d^\Delta p/100)^b |X_i|$. By *(PRGA 5)*, for all but at most $\varepsilon' p^\Delta |V_i|$ vertices v of $V_i' \setminus U$, we have

$$\left|\left\{v' \in V_i : \left|\{x \in X_i^{\text{buf}} : v, v' \in C(x)\}\right| > 24\mu\Delta^2 \left(\tfrac{20}{\mu\zeta d^\Delta}\right)^b p^{2b}|X_i|\right\}\right| \le \varepsilon' p^\Delta |V_i|. \tag{4.4}$$

We choose an ordering $v_1, \dots, v_{|V_i' \setminus U|}$ of $V_i' \setminus U$ which puts the vertices v failing (4.4) first. We choose, for each $j = 1, \dots, |V_i' \setminus U|$, a vertex x_j uniformly at random from the vertices of $X_i^{\text{buf}} \setminus \{x_1, \dots, x_{j-1}\}$ for which v_j is a candidate if this is possible; if not, we say x_j does not exist. As in the previous claim, our aim is to show that a.a.s. all x_j exist, which implies the claim.

Let $\mathscr{H}_0'$ be the empty history, and for each $1 \le j \le |V_i' \setminus U|$, let $\mathscr{H}_j'$ be the history of this process up to and including choosing x_j. We claim that, with positive probability, at each step j we choose from a set of size at least $\frac{1}{2}|C(v_j)| \ge \frac{1}{2}\mu(d^\Delta p/100)^b|X_i|$. This in particular implies that each x_j exists.

We introduce random variables $\tilde{Z}_i^{(j)}$ for $j \in [|V_i' \setminus U|]$ and $\ell \in [j-1]$ as follows. We set $\tilde{Z}_\ell^{(j)} = 1$ if $x_\ell \in C(v_j)$ and $|C(v_\ell) \setminus \{x_1,\ldots,x_{\ell-1}\}| \ge \frac{1}{2}|C(v_\ell)|$ hold, and we set $\tilde{Z}_\ell^{(j)} = 0$ otherwise. Observe that $\sum_{\ell=1}^{j-1} \tilde{Z}_\ell^{(j)}$ is exactly the number of vertices x_k among $x_1, \ldots, x_{j-1}$ which lie in $C(v_j)$ and which were chosen from a set of size at least $\frac{1}{2}|C(v_k)|$. Thus, to prove the claim it suffices to show that with positive probability, for every j we have

$$\sum_{\ell=1}^{j-1} \tilde{Z}_\ell^{(j)} < \tfrac{1}{2}\mu(d^\Delta p/100)^b|X_i|. \tag{4.5}$$

If v_j is a vertex failing (4.4), then $j \le \varepsilon' p^\Delta |V_i| < \frac{1}{4}\mu(d^\Delta p/100)^b|X_i|$, where the second inequality is by choice of ε'. Thus, even if all vertices $x_1,\ldots,x_{j-1}$ happen to be candidates for v_j, we still have (4.5). Next we assume that v_j satisfies (4.4). For each of the at most $\varepsilon' p^\Delta |X_i|$ vertices v_ℓ such that

$$\left|\{x \in X_i^{\text{buf}} : v_j, v_\ell \in C(x)\}\right| > 24\mu\Delta^2\left(\tfrac{20}{\mu\zeta d^\Delta}\right)^b p^{2b}|X_i|,$$

the conditional expectation $\mathbb{E}(\tilde{Z}_\ell^{(j)}|\mathscr{H}_{\ell-1}')$ can be bounded from above by 1, while for the remaining vertices, it is at most

$$\frac{\left|\{x \in X_i^{\text{buf}} : v_j, v_\ell \in C(x)\}\right|}{\frac{1}{2}|C(v_\ell)|} \le \frac{24\mu\Delta^2\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^b p^{2b}}{\frac{1}{2}\mu(d^\Delta p/100)^b} \le 48\Delta^2\left(\tfrac{2000}{\mu\zeta d^{2\Delta}}\right)^\Delta p^b.$$

This yields

$$\sum_{\ell=1}^{j-1} \mathbb{E}\left[\tilde{Z}_\ell^{(j)}\middle|\mathscr{H}_{\ell-1}'\right] \le \varepsilon' p^\Delta |X_i| + |V_i' \setminus U| \cdot 48\Delta^2\left(\tfrac{2000}{\mu\zeta d^{2\Delta}}\right)^\Delta p^b \le \tfrac{1}{4}\mu(d^\Delta p/100)^b|X_i|,$$

where the last inequality is because $|V_i' \setminus U| \le \rho|X_i|$ and by choice of ρ and ε'. As before, we apply Lemma 2.2 for each $j \in [|V_i' \setminus U|]$ with $\delta = 1$ and $(\tilde{Z}^{(j)})_{\ell \in [j-1]}$. We thus may bound the probability that v_j is candidate for more than $\frac{1}{2}|C(v_i)| \ge \frac{1}{2}\mu(d^\Delta p/100)^b|X_i|$ of the vertices $x_1,\ldots,x_{j-1}$ by at most $e^{-\mu(d^\Delta p/100)^b|X_i|/12}$. Taking a union bound over the at most $\rho|X_i|$ choices of j we see that with probability at least $1 - \rho|X_i| \cdot e^{-\mu(d^\Delta p/100)^b|X_i|/12} > 0$ (where the inequality is by choice of p), at each step we choose from a set of size at least $\frac{1}{2}\mu(d^\Delta p/100)^b|X_i|$ as claimed. In particular, we succeed in choosing $|V_i' \setminus U|$ distinct vertices of $X_i^{\text{buf}} \setminus Y$, so as desired we have $|U| \ge |Y|$. □

The final case is to show that $\rho|X_i| < |Y| \le |X_i^{\text{buf}}| - \rho|X_i| = |V_i'| - \rho|V_i|$ and $|U| < |Y|$ is a contradiction. In this case, we have $|V_i' \setminus U| > \rho|V_i|$. By (PRGA 3) there are at most $\rho|X_i|$ vertices of X_i^{buf} with fewer than $\frac{1}{2}(dp)^b|V_i' \setminus U|$ candidates in $V_i' \setminus U$, so in particular there is a vertex of Y with candidates in $V_i' \setminus U$, in contradiction to the definition of U.

This completes the verification of Hall's condition, so we can extend ψ_{RGA} to an embedding ψ of H into G as desired, completing the proof of Lemma 1.25. □

Algorithm 3: Random greedy algorithm for bijumbled graphs

Input : $G \subseteq \Gamma$ and H with partitions satisfying the General Setup; an ordering τ on X^{main}

$t := 0$;

$\psi_0 := \emptyset$;

$Q_0 := \{ x \in V(H) : |I_x| < \frac{1}{2}\mu(d-\varepsilon)^{|J_x|}p^{|J_x|}|V^{\mathrm{main}}(x)| \}$;

repeat

    let $x \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_t)$ be the next vertex in the order τ ;

    if $x \in Q_t$ and $|A_t^{\mathrm{q}}(x) \setminus B_t(x)| < \frac{1}{10}\mu\zeta(dp)^{\pi_t^*(x)}|V(x)|$ **then**

      | halt with failure ;

    end

    choose v uniformly at random in $\begin{cases} A_t^{\mathrm{main}}(x) \setminus B_t(x) & \text{if } x \notin Q_t \\ A_t^{\mathrm{q}}(x) \setminus B_t(x) & \text{if } x \in Q_t \end{cases}$;

    $\psi_{t+1} := \psi_t \cup \{x \to v\}$;

    $Q_{t+1} := Q_t$;

    forall $y \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_{t+1})$ **do**

      if $\left(\left|A_{t+1}^{\mathrm{main}}(y)\right| < \frac{1}{2}\mu(d-\varepsilon')^{\pi_{t+1}^*(y)}p^{\pi_{t+1}^*(y)}|V^{\mathrm{main}}(y)|\right)$ **then**

        | $Q_{t+1} := Q_{t+1} \cup \{y\}$;

      end

    end

    $t := t+1$;

until $\mathrm{Dom}(\psi_t) = X^{\mathrm{main}}$;

$t_{\mathrm{RGAend}} := t$;

4.2 Proof of the bijumbled graphs RGA lemma

The proof of Lemma 4.1 is very similar to the proof of Lemma 3.1 (RGA lemma for random graphs), and indeed we can reuse the auxiliary lemmas of Section 2.4. We again use a random greedy algorithm, Algorithm 3, which we show produces the desired embedding with high probability. The difference from Algorithm 1 is that when at time t we reach a vertex x in the order τ (given again by Lemma 2.28) which is in the queue Q_t, we do not skip it, but instead we embed it into V^{q}. Thus, we always have $t = \tau(x) - 1$. The embedding is done uniformly at random into the set $A_t^{\mathrm{q}}(x) \setminus B_t(x)$. We remind the reader that $B_t(x)$ is the set of bad vertices with respect to ψ_t and Q_t, that is, the vertices v such that the extension $\psi_t \cup \{x \to v\}$ is not a good partial embedding or there is an unembedded neighbour y of x not in Q_t such that $\deg_G\left(v; A_t^{\mathrm{main}}(y)\right) < (d-\varepsilon')p|A_t^{\mathrm{main}}(y)|$ holds (see Definition 2.26).

Observe that there is another important difference from Algorithm 1. Algorithm 3 can fail, and we will have to prove that with high probability it does not. This means proving that sets $A_t^{\mathrm{q}}(x)$ do not get small, or equivalently that sets $C_t^{\mathrm{q}}(x)$ are never substantially covered by $\mathrm{Im}(\psi_t)$. However, there is also an important similarity: provided the algorithm has not yet halted with failure, it maintains the same invariants as Algorithm 1 (see Claim 4.4).

Observe that (in contrast to Algorithm 1) if Algorithm 3 has not yet failed at time t, then all of the first

t vertices in the order τ have been embedded; the vertices are not reordered. Thus $\pi^*_{\tau(x)-1}(x)$ is always equal to $\pi^\tau(x)$, where

$$\pi^\tau(x) := \left|\{y \in N_H(x) : \tau(y) < \tau(x)\}\right| + |J_x|.$$

Much of the work of proving Lemma 4.1 is contained in the auxiliary lemmas in Section 2.4 which we use again here. What remains is to justify that Algorithm 3 with high probability does not halt with failure and does give a good partial embedding with properties (*PRGA* 4) and (*PRGA* 5).

Proof of Lemma 4.1. We require

$$\mu \leq \frac{d^\Delta}{1320\kappa\Delta_{R'}}, \qquad\qquad \rho \leq \frac{\mu^2\zeta^2 d^{2\Delta}}{2000},$$

$$\varepsilon' \leq \frac{\mu\zeta d^\Delta \rho}{1000\kappa\Delta^2\Delta_{R'}4^\Delta}2^{-4/\rho}, \qquad\qquad \varepsilon \leq \min\left((\varepsilon')^2, 2^{-\Delta}\varepsilon', \frac{\rho\mu\zeta d^\Delta}{20\kappa\Delta}, \frac{\varepsilon'}{\kappa\Delta_{R'}}\right),$$

$$p \geq \frac{1000\Delta^4\kappa r_1^2}{\mu\zeta\rho d^\Delta\varepsilon}\left(\frac{\log n}{n}\right)^{1/(2\Delta)} \qquad \text{and} \qquad 2^{-\varepsilon pn/(\kappa r_1)}r_1 n < 1.$$

Recall that we assumed that Γ possesses the three properties $\mathrm{NS}(\varepsilon, r_1, \Delta+1)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, \Delta+1)$ and $\mathrm{LNS}(\varepsilon, r_1, \Delta)$. Let τ be an order on $V(H)$ given by Lemma 2.28. We run Algorithm 3, which maintains the same invariants as Algorithm 1, as shown in the following claim.

Claim 4.4. *The following hold at each time t in the running of Algorithm 3.*

(*INV* 1) ψ_t *is a good partial embedding,*

(*INV* 2) *Either* $|A_t^{\mathrm{main}}(x)| \geq \frac{1}{2}\mu(d-\varepsilon')^{\pi^*_t(x)}p^{\pi^*_t(x)}|V^{\mathrm{main}}(x)|$ *or* $x \in Q_t$.

(*INV* 3) *When we embed x to create ψ_{t+1}, we do so uniformly at random into a set of size at least* $\frac{1}{10}\mu\zeta(dp)^{\pi^*_t(x)}|V(x)|$.

Proof. Algorithm 3 maintains (*INV* 1) and (*INV* 2) by definition. When we embed x to create ψ_{t+1}, either $x \in Q_t$ or not. In the former case, because Algorithm 3 has not failed we have (*INV* 3). In the latter case, in particular x was not added to Q_t at time $t-1$. Using Lemma 2.27 to bound the size of $B_t(x)$, and since $\frac{1}{2}\mu(d-\varepsilon')^b p^b - 20\Delta^2\varepsilon' p^b > \frac{1}{10}\mu\zeta d^b p^b$ by choice of ε' for each $0 \leq b \leq \Delta$, also in this case Algorithm 3 maintains (*INV* 3). □

As in the proof of Lemma 3.1, the conditions for Lemma 2.29 are met, so we conclude that with probability at least $1 - r_1 2^{-n/(\kappa r_1)}$, at each time t when Algorithm 3 is running and for each $i \in [r]$ we have $|Q_t \cap X_i| \leq \rho|X_i| + |X_i^*| \leq 2\rho|X_i|$, where the final inequality is by (*G* 6). In order to show that Algorithm 3 runs successfully, we need to show that the 'halt with failure' line is a.a.s. never reached, i.e. that if $x \in Q_{\tau(x)-1}$ then $\left|A^{\mathrm{q}}_{\tau(x)-1}(x) \setminus B_{\tau(x)-1}(x)\right|$ is at least $\frac{1}{10}\mu\zeta(dp)^{\pi^*_{\tau(x)-1}(x)}|V_i|$. Since we know by Lemma 2.27 that $B_{\tau(x)-1}(x)$ is small and by (*GPE* 3) that $C^{\mathrm{q}}_{\tau(x)-1}(x)$ is large, what we want is to show that $\mathrm{Im}(\psi_{t-1})$ covers only a small fraction of $C^{\mathrm{q}}_{\tau(x)-1}(x)$. We will show this in Claim 4.6. Before this, we argue in Claim 4.5 that in the embedding no 'dense spots' are created, which we need for Claim 4.6.

We now explain how the following Claim 4.5 helps us to establish Claim 4.6. Observe that we embed the vertices $Q_t \cap X_i$ into V_i^{q} in a random procedure which is very similar to the embedding strategy we used in the proof of Claim 4.2 (verification of Hall's condition for small sets of buffer vertices). However, here a complication is that we do not know 'in advance' what $C^{\mathrm{q}}_{\tau(x)-1}(x)$ will be when embedding

earlier vertices to V_i^q. But we do know that it will be contained in the Γ-neighbourhood in V_i of some collection of $\pi^\tau(x) \leq \Delta$ vertices, and that it will be large. Thus, it suffices to prove the following. For each i, each $1 \leq \ell \leq \Delta$ and $v_1, \ldots, v_\ell$ such that $N_\Gamma^*(v_1, \ldots, v_\ell; V_i)$ is large, $\text{Im}(\psi_t)$ never covers much of $N_\Gamma^*(v_1, \ldots, v_\ell; V_i)$. We can prove this statement by using the same analysis as in the proof of Lemma 1.21, given a 'sum condition' (in (*a*) below) similar to (*PRGA* 4). We also show that if Algorithm 3 does not halt with failure then it is likely to have (*PRGA* 4) (which is (*b*) below), since the arguments are similar.

Claim 4.5. *Suppose that for each $i \in [r]$ and each time t we have $|Q_t \cap X_i| \leq 2\rho|X_i|$. Then a.a.s. the following statements hold.*

(*a*) *For any t, any $1 \leq \ell \leq \Delta$, any $i \in [r]$ and any vertices $v_1, \ldots, v_\ell$ such that*

$$\deg_\Gamma(v_1, \ldots, v_\ell; V_i) \geq \varepsilon' p^\ell |V_i|$$

we have

$$\sum_{\substack{y \in X_i^{\text{main}} \cap Q_t: \\ \tau(y) \leq t, J_y = \emptyset}} \frac{\left|U_{\tau(y)-1}(y) \cap N_\Gamma^*(v_1, \ldots, v_\ell; V_i)\right|}{\left|A_{\tau(y)-1}^q(y) \setminus B_{\tau(y)-1}(y)\right|} \leq \frac{\mu \zeta d^\Delta \deg_\Gamma(v_1, \ldots, v_\ell; V_i)}{20}. \tag{4.6}$$

(*b*) *For any t, for any $1 \leq \ell \leq \Delta$, any $i \in [r]$ and any vertices $v_1, \ldots, v_\ell$ such that*

$$\deg_G(v_1, \ldots, v_\ell; V_i^{\text{buf}}) \geq (dp - \varepsilon'p)^\ell |V_i^{\text{buf}}|,$$

we define X_i' to be the set

$$\left\{y \in X_i^{\text{buf}} : \left|U_t^{\text{buf}}(y) \cap N_G^*(v_1, \ldots, v_\ell; V_i^{\text{buf}})\right| > (p + \varepsilon p)^{\pi_t^*(y)} \deg_G(v_1, \ldots, v_\ell; V_i^{\text{buf}})\right\}$$

and we have

$$\sum_{y \in X_i'} \frac{\left|U_t^{\text{buf}}(y) \cap N_G^*(v_1, \ldots, v_\ell; V_i^{\text{buf}})\right|}{\left|C_t^{\text{buf}}(y)\right|} \leq \frac{\deg_G(v_1, \ldots, v_\ell; V_i^{\text{buf}})}{20}. \tag{4.7}$$

The idea of the proof of (4.6) is as follows. The denominators on the left hand side of (4.6) are by (*INV* 3) never much smaller than they 'should' be, so the main task is to show that the numerators do not tend to be too large. To show this, we consider the evolution of $U_t(y) \cap N_\Gamma^*(v_1, \ldots, v_\ell; V_i)$ for some fixed y as t increases. Since $J_y = \emptyset$, at first this set has the size we expect, namely it is all of $N_\Gamma^*(v_1, \ldots, v_\ell; V_i)$. Each time a neighbour of y is embedded, we expect that the set size shrinks by a factor roughly p. If this is the case for each neighbour, the size at $t = \tau(y) - 1$ is roughly a $p^{\pi^{\tau(y)}}$-factor times its original size, which turns out to be a good enough bound for (4.6). If not, there is some first time when we embed a neighbour of y, say the sth neighbour, 'badly', that is, the set size does not shrink by a factor roughly p. We say we fail at step s. At worst, it could be that the set size does not thereafter change, so that it stays roughly a p^{s-1}-factor times its original size. In this case we 'lose' a $p^{\pi^{\tau(y)}-s+1}$ factor. But the LNS property tells us that the probability of failing at step s is less than $p^{\pi^{\tau(y)}-s+1}$. Heuristically, this gets us back the lost factor; to make this rigorous, we apply Lemma 2.2.

Proof of Claim 4.5. We require

$$\rho \leq \frac{\mu^2 \zeta^2 d^{2\Delta}}{2000}, \quad \varepsilon \leq \min\left(2^{-\Delta}\varepsilon', \frac{\rho\mu\zeta d^\Delta}{20\kappa\Delta}\right) \quad \text{and} \quad p \geq \frac{20\Delta^4 \kappa r_1^2}{\rho}\left(\frac{\log n}{n}\right)^{1/\Delta}.$$

We start with (4.6). Since the sum in (4.6) is monotonically increasing in t, and since the upper bound claimed in (4.6) does not depend on t, it suffices to show that the bound holds when Algorithm 3 terminates (with or without failure).

Given $1 \leq \ell \leq \Delta$, let $v_1, \dots, v_\ell$ be vertices of Γ such that the set $U := N_\Gamma^*(v_1, \dots, v_\ell; V_i)$ has size at least $\varepsilon' p^\ell |V_i|$. By (*INV* 3) we have the lower bound

$$\left| A_{\tau(y)-1}^{\mathrm{q}}(y) \setminus B_{\tau(y)-1}(y) \right| \geq \tfrac{1}{10} \mu \zeta (dp)^{\pi^{\tau(y)}} |V_i|$$

for the denominator in each summand of (4.6), and the difficulty is to upper bound the numerator. Consider the running of Algorithm 3. At any time $t \leq \tau(y) - 1$, the set $U_t^{\mathrm{q}}(y) \cap U$ is the Γ-neighbourhood in U of the $\pi_t^*(y)$ embedded vertices from $N_H(y)$, and in a truly random set we would thus expect to find that

$$|U_t(y) \cap U| \leq (p + \varepsilon p)^{\pi_t^*(y)} |U|. \tag{4.8}$$

If this inequality remains true up to $t = \tau(y) - 1$, then since $|J_y| = 0$ by assumption, we have $\pi_{\tau(y)-1}^*(y) = \pi^\tau(y)$, and (4.8) gives an upper bound good enough for the summand $\dfrac{\left| U_{\tau(y)-1}(y) \cap N_\Gamma^*(v_1, \dots, v_\ell; V_i) \right|}{\left| A_{\tau(y)-1}^{\mathrm{q}}(y) \setminus B_{\tau(y)-1}(y) \right|}$. However, it is likely that some vertices will not satisfy (4.8), and we have to estimate their contribution to (4.6).

We say that a vertex $y \in X_i^{\mathrm{main}}$ *fails at step* s if the vertex z is the sth vertex of $N_H(y)$ in τ and y satisfies (4.8) for each $t \leq \tau(z) - 1$ but fails (4.8) at $t = \tau(z)$. Observe that each vertex y in the sum (4.6) satisfies (4.8) before any neighbour of y is embedded because $J_y = \emptyset$. Thus, if y does not fail at any step s with $1 \leq s \leq \Delta$ then it satisfies (4.8) at time $t = \tau(y) - 1$, but if for some $1 \leq s \leq \Delta$ it fails at step s, then $|U_{\tau(y)-1}(y) \cap U| \leq (p + \varepsilon p)^{s-1} |U|$.

Suppose now that y fails at step s. Then the reason is that the vertex z, which is the sth neighbour of y in τ, is embedded to a vertex of Γ with more neighbours in $U_{\tau(z)-1}(y) \cap U$ than allowed by (4.8). Let W be a superset of $U_{\tau(z)-1}(y) \cap U$ of size $(p + \varepsilon p)^{s-1} |U|$. By choice of ε we have $|W| \geq \varepsilon p^{\ell+s-1} |V_i|$. Because y fails at step s we see that z is embedded to a vertex v of Γ with $\deg_\Gamma(v; W) > (p + \varepsilon p)^s |U| = (p + \varepsilon p)|W|$. By LNS$(\varepsilon, r_1, \Delta)$ we know that the number of such vertices v in Γ is at most $\varepsilon p^{2\Delta-s} |V(\Gamma)|/r_1^2$ (since $\ell \leq \Delta$). Therefore, by (*INV* 3), the probability of embedding z (which has at least one unembedded neighbour and hence at most $\Delta - 1$ embedded neighbours) to such a vertex v, conditioning on the history up to but not including embedding z, is at most

$$\frac{\varepsilon p^{2\Delta-s} |V(\Gamma)|/r_1^2}{\mu \zeta (dp)^{\Delta-1} |V_i|/10} \leq \frac{\rho}{2\Delta} p^{\Delta-s+1} \leq \frac{\rho}{2\Delta} p^{\pi^\tau(y)-s+1},$$

where the first inequality is by choice of ε and the second since $\pi^\tau(y) \leq \Delta$.

We can restate this as: the probability that y fails at step $s = \pi^\tau(y) - j + 1$, conditioning on the history up to but not including the embedding of the vertex at step s, is at most $\frac{\rho}{2\Delta} p^j$ for each $1 \leq j \leq \pi^\tau(y)$. It follows that the expected number of vertices $y \in X_i^{\mathrm{main}}$ which fail at step $\pi^\tau(y) - j + 1$ is at most $\frac{\rho}{2\Delta} p^j |X_i|$, and by Lemma 2.2, applied with $\delta = 1$, the probability that more than $\frac{\rho}{\Delta} p^j |X_i|$ vertices $y \in X_i^{\mathrm{main}}$ fail at step $\pi^\tau(y) - j + 1$ is at most $\exp\left(-\frac{\rho}{6\Delta} p^j |X_i|\right)$. Let $\mathcal{E}$ be the event that, for each $1 \leq j \leq \Delta$, at most $\frac{\rho}{\Delta} p^j |X_i|$ vertices y of X_i^{main} fail at step $\pi^\tau(y) - j + 1$. The probability that $\mathcal{E}$ holds is thus at least

$$1 - \Delta \exp\left(-\frac{\rho}{6\Delta} p^\Delta |X_i|\right). \tag{4.9}$$

Suppose $\mathcal{E}$ occurs. Since a vertex y failing at step j satisfies

$$\left|U_{\tau(y)-1}(y)\cap U\right| \le (p+\varepsilon p)^{\pi^\tau(y)-j}|U|$$

we have

$$\sum_{\substack{y\in X_i^{\text{main}}:J_y=\emptyset,\\ |U_{\tau(y)-1}(y)\cap U|>(p+\varepsilon p)^{\pi^\tau(y)}|U|}} \frac{\left|U_{\tau(y)-1}(y)\cap U\right|}{\left|A_{\tau(y)-1}^{\text{q}}(y)\setminus B_{\tau(y)-1}(y)\right|}$$

$$\le \sum_{\substack{y\in X_i^{\text{main}}:\\ y \text{ fails at step } s=\pi^\tau(y)-j+1}} \frac{(p+\varepsilon p)^{\pi^\tau(y)-j}|U|}{\frac{1}{10}\mu\zeta(dp)^{\pi^\tau(y)}|V_i|} \le \sum_{j=1}^{\Delta} \frac{(p+\varepsilon p)^{\pi^\tau(y)-j}|U|\cdot\frac{\rho}{\Delta}p^j|X_i|}{\frac{1}{10}\mu\zeta(dp)^{\pi^\tau(y)}|V_i|}$$

$$\le \frac{20\rho|U|}{\mu\zeta d^\Delta} \le \frac{\mu\zeta d^\Delta|U|}{40}, \quad (4.10)$$

where the last two inequalities are by choice of ε and ρ respectively. On the other hand, because $|X_i^{\text{main}}\cap Q_t| \le 2\rho|X_i|$, we have

$$\sum_{\substack{y\in X_i^{\text{main}}\cap Q_t:\\ |U_{\tau(y)-1}(y)\cap U|\le(p+\varepsilon p)^{\pi^\tau(y)}|U|}} \frac{\left|U_{\tau(y)-1}(y)\cap U\right|}{\left|A_{\tau(y)-1}^{\text{q}}(y)\setminus B_{\tau(y)-1}(y)\right|} \le$$

$$\sum_{y\in X_i^{\text{main}}\cap Q_t} \frac{(p+\varepsilon p)^{\pi^\tau(y)}|U|}{\frac{1}{10}\mu\zeta(dp)^{\pi^\tau(y)}|V_i|} \le \frac{(1+\varepsilon)^\Delta|U|\cdot 2\rho|X_i|}{\frac{1}{10}\mu\zeta d^\Delta|V_i|}$$

$$\le \frac{40\rho|U|}{\mu\zeta d^\Delta} \le \frac{\mu\zeta d^\Delta|U|}{40},$$

and putting these together we conclude (4.6).

It remains to bound the probability that for some vertices $v_1,\dots,v_\ell$ and for some i for which

$$\deg_\Gamma(v_1,\dots,v_\ell;V_i) \ge \varepsilon' p^\ell|V_i|,$$

the inequality (4.6) is violated, in other words that $\mathcal{E}$ fails for $v_1,\dots,v_\ell$ and i. There are $r \le r_1$ choices of i and at most $n+n^2+\cdots+n^\Delta \le \Delta n^\Delta$ choices of $v_1,\dots,v_\ell$, so by the union bound and (4.9) the probability of such a bad event occurring is at most

$$r_1\Delta n^\Delta \cdot 2\Delta\exp\left(-\frac{\rho}{6\Delta\kappa r_1}p^\Delta n\right)$$

which tends to zero as n tends to infinity by choice of p.

We now come to (4.7). This inequality looks very much like (4.10) above, and its proof is almost identical to the proof of (4.10). The differences are that we consider G-neighbourhoods in V_i^{buf} not Γ-neighbourhoods in V_i and that we have a lower bound on $C_t^{\text{buf}}(x)$ from $(GPE\,3)$ rather than on $|A_t^{\text{q}}(x)\setminus B_t(x)|$ from $(INV\,3)$. Since this lower bound is larger, the same constant choices work. We omit the details.   □

We can now use Lemma 2.2 to complete the proof that the sets $C_t^q(x)$ do not get covered by $\mathrm{Im}(\psi_t)$, much as in Claim 4.2 (verification of Hall's condition for small sets of buffer vertices).

Claim 4.6. *A.a.s. for each $x \in X^{\mathrm{main}}$, at each time $t \leq \tau(x) - 1$ and before the termination of Algorithm 3, we have $\left| C_t^q(x) \cap \mathrm{Im}(\psi_t) \right| < \frac{1}{2} \left| C_t^q(x) \right|$.*

Proof. We require

$$\varepsilon \leq \varepsilon' \leq 2^{-2\Delta}, \qquad \rho \leq \frac{\mu \zeta d^{\Delta}}{30} \qquad \text{and} \qquad p \geq \frac{1000 r_1}{\mu \zeta d^{\Delta}} \left(\frac{\log n}{n} \right)^{1/\Delta}.$$

Suppose that the conclusions of Claim 4.5 hold (in particular (4.6) holds). Given $i \in [r]$ and $x \in X_i^{\mathrm{main}}$, if $t \leq \tau(x) - 1$ and Algorithm 3 has not terminated before time t, then by $(GPE\,3)$ and $(G\,5)$ we have $\left| C_t^q(x) \right| \geq (1 - \varepsilon') \mu \zeta (dp - \varepsilon' p)^{\pi_t^*(x)} |V_i|$. Since $C_t^q(x) \subseteq U_t(x) = N_{\Gamma}^*(v_1, \ldots, v_\ell; V_i)$ for some vertices $v_1, \ldots, v_\ell$ with $\ell = \pi_t^*(x) \leq \Delta$, we conclude by (4.6) that

$$\sum_{\substack{y \in X_i^{\mathrm{main}} \cap Q_t: \\ \tau(y) \leq t, J_y = \emptyset}} \frac{\left| U_{\tau(y)-1}^q(y) \cap C_t^q(x) \right|}{\left| A_{\tau(y)-1}^q(y) \setminus B_{\tau(y)-1}(y) \right|} \leq \frac{\mu \zeta d^{\Delta} \left| U_t(x) \right|}{20}.$$

Observe that the summand in the above inequality is an upper bound for the probability that y is embedded to $C_t^q(x)$ for $y \in Q_t$, conditioning on the history up to but not including the embedding of y. Since the probability that y is embedded to $C_t^q(x)$ is zero if y is not in Q_t, we are in a position to apply Lemma 2.2. This lemma, with $\delta = 1$, tells us that the probability that more than $\frac{1}{10} \mu \zeta d^{\Delta} |U_t(x)|$ vertices y of X_i^{main} with $J_y = \emptyset$ are embedded to $C_t^q(x)$ is at most $\exp \left(- \frac{1}{60} \mu \zeta d^{\Delta} |U_t(x)| \right)$. If this bad event does not occur, then by $(GPE\,2)$, $(GPE\,3)$ and $(G\,5)$, and because the number of vertices in X_i^{main} with $J_y \neq \emptyset$ is by $(G\,6)$ at most $\rho p^{\Delta} |X_i|$, we have the desired statement:

$$\left| C_t^q(x) \cap \mathrm{Im}(\psi_t) \right| < \frac{1}{10} \mu \zeta d^{\Delta} \left| U_t(x) \right| + \rho p^{\Delta} |X_i| \leq \frac{1}{2} (1 - \varepsilon') \mu \zeta (dp - \varepsilon' p)^{\pi_t^*(x)} |V_i|.$$

The probability that the conclusions of Claim 4.5 fail to hold, or that any of the above bad events occur, is at most $o(1) + r_1 n^2 \cdot \exp \left(- \frac{1}{60 r_1} \mu \zeta d^{\Delta} (p - \varepsilon' p)^{\Delta} n \right)$, which tends to zero as n tends to infinity by choice of p, completing the proof. □

If the conclusions of Lemma 2.29, Claim 4.5 and Claim 4.6 hold (which we think of as being good events), then by $(GPE\,3)$ and Lemma 2.27 we have

$$\left| A_{\tau(x)-1}^q(x) \setminus B_{\tau(x)-1}(x) \right| \geq \frac{1}{2} \mu \zeta (dp - \varepsilon' p)^{\pi^{\tau}(x)} |V_i| - 20 \Delta^2 \varepsilon' p^{\pi^{\tau}(x)} |V_i|$$

and the right hand side is by choice of ε' at least $\frac{1}{10} \mu \zeta (dp)^{\pi^{\tau}(x)} |V_i|$. In other words, the 'halt with failure' line of Algorithm 3 is never reached. Since each of the three good events a.a.s. occurs, the algorithm a.a.s. completes successfully. Now $(PRGA\,1)$ is guaranteed by successful completion of Algorithm 3, and $(PRGA\,3)$ is as $(RGA\,4)$ in the proof of Lemma 3.1 implied by the good event of Lemma 2.29, while $(PRGA\,4)$ is given by the good event of Claim 4.5, specifically taking (4.7) with $t = t_{\mathrm{RGAend}}$ and vertices $\psi_t \left(N_H(x) \right)$ for each $x \in X_i^{\mathrm{buf}}$.

Next we establish that *(PRGA 2)* a.a.s. holds. We apply Lemma 2.30 with $B = \Delta + 1$ and $T = \tau(z)$ where z is the last vertex in $N(X^{\text{buf}})$. Since $N(X^{\text{buf}})$ forms the initial segment of τ by Lemma 2.28(*a*), we see that all vertices y embedded up to time T have $\pi^\tau(y) \le \Delta - 1 \le B - 2$, and by identical logic as in Claim 3.6 none of these vertices enter the queue. Therefore, by *(GPE 3)*, Lemma 2.27 and choice of ε', each y from $N(X^{\text{buf}})$ is embedded uniformly at random to a set of size at least $\frac{1}{10}(dp)^{\pi^\tau(y)}|V(y)|$. Finally, by *(H 3)* we have $|N(X^{\text{buf}}) \cap X_i| \le 4\kappa\Delta_{R'}\mu|X_i|$ for each $i \in [r]$. This justifies that the conditions of Lemma 2.30 are met, so we conclude that a.a.s. the following event $\mathcal{E}_{\text{L2.30}}$ occurs. For each $v \in V_i$ and j such that $ij \in E(R')$ we have $\left|N_G(v; V_j^{\text{main}}) \setminus \text{Im}(\psi_T)\right| \ge \frac{1}{2}\deg_G(v; V_j^{\text{main}})$.

We are now in a position to apply Lemma 2.31, again with $B = \Delta + 1$. Observe that this time, if $\mathcal{E}_{\text{L2.30}}$ occurs, it applies to any $x \in X_i^{\text{buf}}$ and $v \in V_i$, whether or not x is in a clique buffer. The deduction that a.a.s. the embedding $\psi_{t_{\text{RGAend}}}$ has *(PRGA 2)* follows exactly as in the proof of Lemma 3.1 (RGA lemma), and we do not repeat it.

It remains to establish *(PRGA 5)*. For this we require the following claim, which shows that if vertices v and v' are common candidates for too many $x \in X_i^{\text{buf}}$ for *(PRGA 5)*, then the reason is that they have an exceptionally large common Γ-neighbourhood in some cluster.

Claim 4.7. *Asymptotically almost surely at the termination of Algorithm 3 the following holds. For each i and pair of vertices $v, v' \in V_i$ such that $\deg_\Gamma(v, v'; V_j) \le (p + \varepsilon p)^2|V_j|$ whenever $ij \in R'$, if b is such that X_i^{buf} is a degree-b buffer, then we have*

$$\left|\{x \in X_i^{\text{buf}} : v, v' \in C(x)\}\right| \le 24\mu\Delta^2 \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^b p^{2b}|X_i|. \tag{4.11}$$

To prove this we show that for each $x \in X_i^{\text{buf}}$, however, previous vertices are embedded, it is not too likely that $N_H(x)$ is embedded to $N_\Gamma^*(v, v')$, and apply Lemma 2.2 and the union bound to deduce (4.11) for all i, v, and v'. In turn, to prove the desired upper bound on the probability of embedding $N_H(x)$ to $N_\Gamma^*(v, v')$, we analyse the embedding of the vertices $N_H(x) = \{y_1, \dots, y_b\}$ one by one. We would expect that in each case roughly a p^2-fraction of $U(y_i)$ is contained in $N_\Gamma^*(v, v')$, and if this is the case at each step we obtain the desired upper bound. If not, the reason is that a previously embedded vertex—which must be one of the y_j since we have not yet embedded any other neighbouring vertices—was 'badly embedded'. Using the LNS property, we show this is an unlikely event and again obtain the desired upper bound. In fact, this is a slight oversimplification: we have to separate the cases that exactly one vertex is embedded badly (which implies we still have only at worst a probability p of embedding future vertices of $N_H(x)$ to $N_\Gamma^*(v, v')$) or that more than one vertex is embedded badly, in which case we might have probability 1 of embedding future vertices of $N_H(x)$ to $N_\Gamma^*(v, v')$, but this is counterbalanced by the unlikelihood of embedding two vertices badly.

Proof of Claim 4.7. We require

$$\varepsilon \le \min\left(2^{-\Delta}, \frac{\mu\zeta d^\Delta}{10\Delta\kappa}\right) \quad \text{and} \quad p \ge 20\mu^{-1}\kappa r_1\left(\frac{\log n}{n}\right)^{1/(2\Delta)}.$$

Let $v, v' \in V_i$ be such that $\deg_\Gamma(v, v'; V_j) \le (p + \varepsilon p)^2|V_j|$ for each j such that $ij \in R'$. Suppose that X_i^{buf} is a degree-b buffer and $b \in [\Delta]$.

Let $T = \tau(x) - |N_H(x)| - 1$ be the time immediately before the first vertex of $N_H(x)$ is embedded. Since any two vertices of X_i^{buf} are at distance at least ten in H by *(H 2)*, it is enough to show

($*$) for any one vertex $x \in X_i^{\text{buf}}$, by considering the embedding of $N_H(x)$, that the probability of $v, v' \in C(x)$, conditioning on the history up to time T, is at most $3\Delta^2 \big(20\mu^{-1}\zeta^{-1}d^{-\Delta}\big)^b p^{2b}$.

If we show this, then, since the vertices of $N_H(x)$ are embedded consecutively, by Lemma 2.2, applied with $\delta = 1$ and using $|X_i^{\text{buf}}| = 4\mu|X_i|$, we see that (4.11) holds with probability at least $1 - \exp\big(-4\mu\Delta^2\big(20\mu^{-1}\zeta^{-1}d^{-\Delta}\big)^b p^{2b}|X_i|\big)$, and taking a union bound over the at most n^2 choices of v, v' we conclude that by choice of p the conclusion of the claim a.a.s. holds as desired.

We show that the probability of $v, v' \in C(x)$ is at most $3\Delta^2 \big(20\mu^{-1}\zeta^{-1}d^{-\Delta}\big)^b p^{2b}$, conditioning on any ψ_T. Observe that $v, v' \in C(x)$ can occur only if $N_H(x)$ is embedded into $N_\Gamma^*(v, v')$. If $y \in N_H(x) \cap X_j$, then the probability that y is embedded into $N_\Gamma^*(v, v')$ is, using (*INV 3*), at most

$$\frac{\big|U_{\tau(y)-1}(y) \cap N_\Gamma^*(v, v')\big|}{\frac{1}{10}\mu\zeta(dp)^{\pi^\tau(y)}|V_j|}.$$

We would expect that the numerator is at most $(p + \varepsilon p)^{2+\pi^\tau(y)}|V(y)|$ for all $y \in N_H(x)$. Of course, this does not always happen, since it may be the case that a neighbour z of y in H with $\tau(z) < \tau(y)$ was embedded *badly*, that is, to a vertex with more than

$$(p + \varepsilon p)\big|U_{\tau(z)-1}(y) \cap N_\Gamma^*(v, v')\big| \tag{4.12}$$

neighbours in $U_{\tau(z)-1}(y) \cap N_\Gamma^*(v, v')$. Since τ satisfies the condition of Lemma 2.28(*a*), any such z is in $N(X^{\text{buf}})$. Since x is at distance at most 2 from z in H, and at distance at least 5 from any other vertex of X^{buf} by (*H 4*), we see that $z \in N_H(x)$. Thus, if exactly s many neighbours z of y with $\tau(z) < \tau(y)$ are embedded badly then we see that with (4.12) the following inequality

$$\big|U_{\tau(y)-1}(y) \cap N_\Gamma^*(v, v')\big| \le (p + \varepsilon p)^{\pi^\tau(y)-s}|N_\Gamma^*(v, v')| \le (p + \varepsilon p)^{\pi^\tau(y)+2-s}|V(y)| \tag{4.13}$$

holds. We now show how this implies ($*$). We separate three cases.

First, no vertices in $N_H(x)$ are embedded badly. Then the probability that all b vertices in $N_H(x)$ are embedded into $N_\Gamma^*(v, v')$, conditioning on ψ_T, is at most $\big(20\mu^{-1}\zeta^{-1}d^{-\Delta}p^2\big)^b$ by choice of ε.

Second, exactly one vertex in $N_H(x)$ is embedded badly. We let the vertices of $N_H(x)$ be $y_1, \dots, y_b$ in order of τ. Let us suppose that y_ℓ is the vertex embedded badly. Then the probability that the first $\ell - 1$ vertices of $N_H(x)$ are embedded into $N_\Gamma^*(v, v')$, conditioning on ψ_T, is at most $\big(20\mu^{-1}\zeta^{-1}d^{-\Delta}p^2\big)^{\ell-1}$ by choice of ε (since these are not badly embedded and conditional probabilities multiply).

We now estimate the probability of y_ℓ being embedded badly, conditioning on ψ_T and on the embeddings of the previous $\ell - 1$ vertices not being bad. Observe that, since there has been no previous bad embedding, we have for each $\ell' > \ell$

$$\big|U_{\tau(y_\ell)-1}(y_{\ell'}) \cap N_\Gamma^*(v, v')\big| \le (p + \varepsilon p)^{\pi_{\tau(y_\ell)-1}(y_{\ell'})+2}|V(y_{\ell'})|,$$

and since $\pi_{\tau(y_\ell)-1}(y_{\ell'}) \le \ell - 1 \le \Delta - 2$, the right hand side is, by choice of ε, at least $\varepsilon p^\Delta n/r_1$. By $\text{LNS}(\varepsilon, r_1, \Delta)$ the number of vertices in Γ with more than $(p + \varepsilon p)^{\pi_{\tau(y_\ell)-1}(y_{\ell'})+3}|V(y_{\ell'})|$ neighbours into $U_{\tau(y_\ell)-1}(y_{\ell'}) \cap N_\Gamma^*(v, v')$ is at most $\varepsilon p^{2\Delta-1}n/r_1$. Since $\pi^\tau(y_\ell) \le \ell - 1$, by (*INV 3*) the probability that y_ℓ is embedded to such a bad vertex is at most

$$\frac{\Delta\varepsilon p^{2\Delta-1}n/r_1}{\mu\zeta(dp)^{\ell-1}|V(y_\ell)|/10} \le p^{2\Delta-\ell}$$

where the inequality is by choice of ε.

Now, the probability that the last $b - \ell$ vertices are embedded into $N_\Gamma^*(v, v')$, conditioning on the previous embeddings, is at most $\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}p\right)^{b-\ell}$ by choice of ε. Indeed, for each vertex $\ell' > \ell$ we have by (4.13) that

$$\left|U_{\tau(y_{\ell'})-1}(y_{\ell'}) \cap N_\Gamma^*(v, v')\right| \leq (p + \varepsilon p)^{\pi^\tau(y_{\ell'})+1}|V(y_{\ell'})|$$

since we condition on exactly one neighbour of $y_{\ell'}$ being embedded badly. This gives a conditional probability of at most $20\mu^{-1}\zeta^{-1}d^{-\Delta}p$ that $v_{\ell'}$ is embedded to $N_\Gamma^*(v, v')$. The conditional probabilities multiply, and taking the union bound over the choices of ℓ, we see that the probability of this case occurring and all vertices of $N_H(x)$ being embedded to $N_\Gamma^*(v, v')$, conditioning on ψ_T, is at most

$$\Delta \cdot \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}p^2\right)^{\ell-1} \cdot p^{2\Delta-\ell} \cdot \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}p\right)^{b-\ell} \leq$$
$$\Delta\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^{b-1}p^{2\Delta+b-2} \leq \Delta^2\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^b p^{2b},$$

where the last inequality follows since $\Delta \geq 2$ and $b \leq \Delta$.

In the third case, at least two vertices in $N_H(x)$ are embedded badly. Suppose that the first two badly embedded vertices are the jth and kth vertices. Then the same logic as above, in particular the inequality (4.13), tells us that the probability that this case occurs and all vertices of $N_H(x)$ are embedded to $N_\Gamma^*(v, v')$ is at most

$$\Delta^2 \cdot \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}p^2\right)^{j-1} \cdot p^{2\Delta-j} \cdot \left(20\mu^{-1}\zeta^{-1}d^{-\Delta}p\right)^{k-j-1} \cdot p^{2\Delta-k} \leq$$
$$\Delta^2\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^{b-2}p^{4\Delta-3}.$$

Putting these three cases together, and using the fact that $4\Delta - 3 \geq 2b$ since $\Delta \geq 2$ and $b \leq \Delta$, we conclude that the probability that $v, v' \in C(x)$ is at most $3\Delta^2\left(20\mu^{-1}\zeta^{-1}d^{-\Delta}\right)^b p^{2b}$ as desired. □

Suppose now that the good event of Claim 4.7 holds. For each $i \in [r]$, there may be exceptional vertices $v \in V_i$ which have more than $(p + \varepsilon p)|V_j|$ neighbours in V_j for some j such that $ij \in R'$, but by $\mathrm{NS}(\varepsilon, r_1, \Delta+1)$ there are at most $\Delta_{R'}\varepsilon p^\Delta n/r_1 \leq \varepsilon' p^\Delta|V_i|$ such vertices. If v is not exceptional—that is, it has at most $(p + \varepsilon)|V_j|$ neighbours in each V_j with $ij \in R'$—then again there may be vertices $v' \in V_i$ which are exceptional for v, i.e. vertices v' with more than $(p + \varepsilon p)^2|V_j|$ common neighbours with v in V_j for some j such that $ij \in R'$. But again by $\mathrm{NS}(\varepsilon, r_1, \Delta+1)$ there are at most $\Delta_{R'}\varepsilon p^\Delta n/r_1 \leq \varepsilon' p^\Delta|V_i|$ such. Because the good event of Claim 4.7 holds, for non-exceptional pairs of vertices v, v' we have the bound given in (*PRGA* 5). Because the good event of Claim 4.7 holds a.a.s., property (*PRGA* 5) is a.a.s. obtained. This completes the proof of Lemma 4.1. □

We note that it is Claim 4.5 in this proof which is most responsible for our bijumbledness requirement on Γ in the blow-up lemma for bijumbled graphs, Lemma 1.25. Although the proof of this claim may seem quite wasteful—we assume that if one neighbour is embedded badly then potentially all future neighbours can be embedded badly without any further penalty, which seems unreasonable—we were not able to make it work with any weaker condition than $\mathrm{LNS}(\varepsilon, r_1, \Delta)$ for general graphs H. Our analysis is in some sense tight for vertices y which fail at step $\pi^\tau(y)$, when the final neighbour z of y coming before y in τ is embedded. What still seems unreasonable is that all the vertices in Γ which have exceptionally high

degree to $U_{\tau(z)-1}(y) \cap U$ turn out to be in $A_{\tau(z)-1}(z)$, which is the worst case we are effectively using in our proof. We expect that a more careful analysis, possibly involving some modification to Algorithm 3, would allow one to improve on this and thus improve on the bijumbledness requirement of Lemma 1.25.

Chapter 5

Improved bounds for degenerate graphs

5.1 The RGA lemma and the proof of the blow-up lemma

In this section we prove Lemma 1.23. We have already seen most of the ideas in the proofs of Lemmas 1.21 and 1.25. We use the same General Setup, and we continue to obtain it using Lemma 2.22. However, rather than defining an order τ putting the buffer vertices X^{buf} at the end and their neighbours at the front, as we did in the proofs of Lemmas 3.1 and 4.1, we are supplied with an order τ which we will modify only in that we move X^{buf} to the end of the order. In particular, the neighbours of buffer vertices need not appear early in τ. The reason for this is that moving vertices in the order could result in substantially increasing $\pi^{\tau}(x)$ for some vertices x. Moving the vertices of X^{buf} to the end of the order can destroy boundedness of the order, but it turns out that the only vertices which cause boundedness to fail are those in X^{buf}, and in the proof the potentially failing conditions are not needed for these vertices. The following definition, which extends (D, p, m)-boundedness, accommodates this.

Definition 5.1 ((D, p, m)-bounded order)**.** Let H be a graph given with buffer sets $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ and a restriction pair $\mathcal{I} = \{I_x\}_{x \in V(H)}$ and $\mathcal{J} = \{J_x\}_{x \in V(H)}$. Let $\tilde{X} = \bigcup_{i \in [r]} \tilde{X}_i$. Let τ be an ordering of $V(H)$ and $X^e, X^{\mathrm{buf}} \subseteq V(H)$. Then τ is a (D, p, m)-*bounded order* for H, $\tilde{\mathcal{X}}$, $\mathcal{I}$ and $\mathcal{J}$ with *exceptional set* X^e and *ignored set* X^{buf} if the following conditions are satisfied for each $x \in V(H)$.

(*ORD* 1') Define

$$D_x := \begin{cases} D - 2 & \text{if there is } yz \in E(H) \text{ with } y, z \in N_H(x) \text{ and } \tau(y), \tau(z) > \tau(x) \\ D - 1 & \text{else if there is } y \in N_H(x) \text{ with } \tau(y) > \tau(x) \\ D & \text{otherwise}. \end{cases}$$

We have $\pi^{\tau}(x) \leq D_x$, and if $x \in N(\tilde{X}) \setminus X^{\mathrm{buf}}$ even $\pi^{\tau}(x) \leq D_x - 1$. Finally, if $x \in \tilde{X}$ we have $\deg(x) \leq D$.

(*ORD* 2') One of the following holds:
   - $x \in X^e \cup X^{\mathrm{buf}}$,
   - $\pi^{\tau}(x) \leq \frac{1}{2}D$,

- x is not image restricted and every neighbour y of x with $\tau(y) < \tau(x)$ satisfies $\tau(x) - \tau(y) \leq p^{\pi^\tau(x)}m$.

$(ORD\,3')$ If $x \in N(\tilde{X}) \setminus X^{\text{buf}}$ then all but at most $D - 1 - \max_{z \notin X^e} \pi^\tau(z)$ neighbours y of x with $\tau(y) < \tau(x)$ satisfy $\tau(x) - \tau(y) \leq p^D m$.

Observe that the only difference between Definitions 1.22 and 5.1 is the introduction of the ignored set X^{buf}, and conditions $(ORD\,2')$ and $(ORD\,3')$, and part of $(ORD\,1')$, are changed so that they are trivially satisfied by vertices of the ignored set X^{buf}. If the ignored set is empty, then Definition 1.22 is recovered.

We will use a modification of Algorithm 3 to perform this embedding. The modification consists of handling the exceptional vertices X^e differently, which allows us to deal with a few vertices x with $\pi^\tau(x)$ significantly larger than normal; such vertices appear in applications. We can show that this algorithm succeeds in embedding all the vertices of X^{main}, in the order τ, and that it returns a good partial embedding with the properties detailed in the lemma below with positive probability. Since we will have $X^{\text{main}} = V(H) \setminus X^{\text{buf}}$, we can complete the proof using Lemma 3.4 to embed the buffer vertices in much the same way as in the proof of Lemma 1.21.

Note that property LCON, which appears in the following lemma, plays the same rôle as LNS in the proof of Lemma 4.1. It is easier to work with, and allows for linearly many image restrictions (but it is not in general true in bijumbled graphs).

Lemma 5.2 (degenerate RGA lemma). *We assume the General Setup. Suppose that an exceptional set X^e with $|X^e| \leq \frac{1}{2}\varepsilon p^{\max_{x \in X^e} \pi^\tau(x)} n/r_1$ is given. Suppose that $D \geq 1$, and τ is a $\left(D, p, \frac{1}{2}\varepsilon n/r_1\right)$-bounded order for H, $\tilde{X}$, $\mathfrak{I}$ and $\mathfrak{J}$ with the exceptional set X^e and ignored set X^{buf}. Suppose furthermore that the vertices X^{buf} form the final segment of τ. Suppose that Γ has properties $\text{NS}(\varepsilon, r_1, D)$, $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$ and $\text{LCON}(\varepsilon, r_1, D)$. Then there is a good partial embedding ψ_{RGA} of H into G with the following properties for each $i \in [r]$. Let b be such that X_i^{buf} is a degree-b buffer.*

$(DRGA\,1)$ *Every vertex in X_i^{main} is embedded to $V_i^{\text{main}} \cup V_i^{\text{q}} \cup V_i^{\text{c}}$ by ψ_{RGA}, and no vertex in X_i^{buf} is embedded.*

$(DRGA\,2)$ *For every set $W \subseteq V_i$ of size at least $\rho|V_i|$, there are at most $\rho|X_i|$ vertices in X_i^{buf} with fewer than $(dp)^b|W|/2$ candidates in W.*

$(DRGA\,3)$ *Every vertex in V_i is a candidate for at least*

$$4^{-10\Delta^3} 2^{-1000D^2 \mu^{-1} \zeta^{-1} d^{-D}} d^{2D^2} \Delta^{-3} \mu p^b |X_i|$$

vertices of X_i^{buf}.

We will prove this lemma in Section 5.2. Assuming Lemma 5.2, we are in a position to prove Lemma 1.23. This amounts to a verification that the conditions of Lemma 1.23 suffice to apply Lemmas 2.17, 2.22, 3.4 and 5.2.

Proof of Lemma 1.23. First we choose constants as follows. Given Δ, $\Delta_{R'}^{\text{BL}}$, Δ_J and D integers, α^{BL}, ζ^{BL} and $d > 0$, and $\kappa^{\text{BL}} > 1$, we set $\vartheta = 0$, $\Delta_{R'} = 8(\Delta + \Delta_J)^{10}\Delta_{R'}^{\text{BL}}$, $\alpha = \frac{1}{2}\alpha^{\text{BL}}$, $\zeta = \frac{1}{2}\zeta^{\text{BL}}$ and $\kappa = 2\kappa^{\text{BL}}$. We

now choose

$$\mu = \frac{\alpha}{20000\kappa\Delta_{R'}^4\Delta^{10}(\Delta+2)}, \quad \rho = \frac{\mu^2\zeta^2d^{3D^2}}{10000\kappa(\Delta+1)}4^{-2000\Delta^3\mu^{-1}\zeta^{-1}d^{-D}}$$

$$\text{and} \quad \varepsilon' = \frac{\mu\zeta d^{\Delta+1}}{1000\Delta^4\kappa4^\Delta}2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}.$$

Now for input D and d, ε', Lemma 2.17 returns constants $(\varepsilon_{a,b})$ and $\varepsilon_{L(b)} > 0$. We set

$$\varepsilon = \min\left(\varepsilon_{L(b)}, \frac{d\varepsilon'}{\kappa}\right).$$

We let $\varepsilon^{BL} = \frac{1}{16}(\Delta+\Delta_J)^{-10}\varepsilon$ and $\rho^{BL} = \frac{1}{16}(\Delta+\Delta_J)^{-10}\rho$. Now Lemma 1.23 returns ε^{BL} and ρ^{BL}. Given r_1^{BL} we let $r_1 = 8(\Delta+\Delta_J)^{10}r_1^{BL}$. We choose C sufficiently large for Lemma 2.17 with input D, d, ε', r_1 and ρ, and such that

$$C \geq \frac{10^8 \cdot 4^{10\Delta^3}(\Delta+\Delta_J)^4\kappa r_1^2}{\varepsilon^2\mu\zeta d^{\Delta+1}}2^{280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}.$$

Given $p \geq C\left(\frac{\log n}{n}\right)^{1/D}$, Lemma 2.17 states that a.a.s. $\Gamma = G(n, p)$ has the properties $NS(\varepsilon, r_1, D)$, $RI(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$, $CON(\rho, r_1, D)$ and $LCON(\varepsilon, r_1, D)$ respectively. From now on we will assume Γ is an n-vertex graph which satisfies these four properties.

Given a graph R_{BL} on $r_{BL} \leq r_1^{BL}$ vertices, and a spanning subgraph R'_{BL} with $\Delta(R'_{BL}) \leq \Delta_{R'}^{BL}$, and graphs H and $G \subseteq \Gamma$ with vertex partitions $\mathcal{X}^{BL}$ and $\mathcal{V}^{BL}$, families of image restrictions $\mathcal{J}^{BL}$ and of image restricting vertices $\mathcal{J}$, a family of potential buffer vertices $\tilde{\mathcal{X}}^{BL}$, and an exceptional set X^e, suppose that the conditions of Lemma 1.23 are satisfied. Then Lemma 2.22 gives a graph R on $r \leq r_1$ vertices, a spanning subgraph R' with $\Delta(R') \leq \Delta_{R'}$, and κ-balanced size-compatible partitions $\mathcal{X}$ and $\mathcal{V}$ of H and G respectively, each part having size at least $n/(\kappa r_1)$, together with a family $\tilde{\mathcal{X}}$ of potential buffer vertices and $\mathcal{J}$ of image restrictions, subsets X_i^{buf} of $\tilde{X}_i$ for each $i \in [r]$, and partitions $V_i = V_i^{main} \dot\cup V_i^q \dot\cup V_i^c \dot\cup V_i^{buf}$ for each $i \in [r]$ which satisfy the General Setup.

We now modify the provided $(D, p, \varepsilon^{BL}n/r_1^{BL})$-bounded order τ^{BL} on $V(H)$ by moving the set X^{buf} to the end of the order, obtaining a new order τ. We claim τ is $(D, p, \varepsilon^{BL}n/r_1^{BL})$-bounded for H, $\tilde{\mathcal{X}}^{BL}$, $\mathcal{J}^{BL}$ and $\mathcal{J}$ with exceptional set X^e and ignored set X^{buf}. To see this, observe that if $x \notin X^{buf}$ then $\pi^\tau(x) \leq \pi^{\tau^{BL}}(x)$. It is easy to check from this that vertices not in X^{buf} satisfy all three conditions. If $x \in X^{buf}$ then $\pi^\tau(x) = \deg(x) \leq D$ by (ORD 1), since $X^{buf} \subseteq \tilde{X}$, and since X^{buf} is an independent set by (H 4) it follows that vertices of X^{buf} satisfy (ORD 1'). Finally vertices in X^{buf} trivially satisfy (ORD 2') and (ORD 3').

Since $\varepsilon^{BL}n/r_1^{BL} = \frac{1}{2}\varepsilon n/r_1$, the order τ is also $(D, p, \frac{1}{2}\varepsilon n/r_1)$-bounded for H, $\tilde{\mathcal{X}}$, $\mathcal{J}$ and $\mathcal{J}$ with exceptional set X^e and ignored set X^{buf}. To see this observe that the set $\mathcal{J}$ and the specific partition $\tilde{\mathcal{X}}$ of $\tilde{X}$ do not play a rôle in Definition 5.1 (the former comes together with the set $\mathcal{J}$, which does play a rôle, and the latter is important only in that it defines $\tilde{X}$).

We let $X^{main} = V(H) \setminus X^{buf}$. We now begin the embedding of H into G. By Lemma 5.2, there is a good partial embedding ψ_{RGA} with the properties stated in that lemma. By (GPE 3), condition (BUF 1) of Lemma 3.4 is satisfied, while condition (BUF 2) holds by (DRGA 2) and (BUF 3) follows from (DRGA 3) with

$$\delta = 4^{-10\Delta^3}2^{-1000D^2\mu^{-1}\zeta^{-1}d^{-D}}d^{2D^2}\Delta^{-3}\mu.$$

Algorithm 4: Random greedy algorithm for degenerate graphs

Input : $G \subseteq \Gamma$ and H with partitions satisfying the General Setup; an ordering τ on X^{main}

$t := 0$;

$\psi_0 := \emptyset$;

$Q_0 := \{x \in V(H) : |I_x| < \frac{1}{2}\mu(d-\varepsilon)^{|J_x|}p^{|J_x|}|V^{\mathrm{main}}(x)|\}$;

repeat

    let $x \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_t)$ be the next vertex in the order τ ;

    if $x \in Q_t \setminus X^e$ and $|A_t^{\mathrm{q}}(x) \setminus B_t(x)| < \frac{1}{10}\mu\zeta(dp)^{\pi_t^*(x)}|V(x)|$ **then**

        | halt with failure ;

    end

    if $x \in Q_t \cap X^e$ and $|A_t^{\mathrm{c}}(x) \setminus B_t(x)| < \frac{1}{10}\mu\zeta(dp)^{\pi_t^*(x)}|V(x)|$ **then**

        | halt with failure ;

    end

    choose v uniformly at random in $\begin{cases} A_t^{\mathrm{main}}(x) \setminus B_t(x) & \text{if } x \notin (Q_t \cup X^e)\,, \\ A_t^{\mathrm{q}}(x) \setminus B_t(x) & \text{if } x \in Q_t \setminus X^e\,, \\ A_t^{\mathrm{c}}(x) \setminus B_t(x) & \text{if } x \in X^e\,; \end{cases}$

    $\psi_{t+1} := \psi_t \cup \{x \to v\}$;

    $Q_{t+1} := Q_t$;

    forall $y \in X^{\mathrm{main}} \setminus \mathrm{Dom}(\psi_{t+1})$ **do**

        if $\left(|A_{t+1}^{\mathrm{main}}(y)| < \frac{1}{2}\mu(d-\varepsilon')^{\pi_{t+1}^*(y)}p^{\pi_{t+1}^*(y)}|V^{\mathrm{main}}(y)|\right)$ **then**

            | $Q_{t+1} := Q_{t+1} \cup \{y\}$;

        end

    end

    $t := t+1$;

until $\mathrm{Dom}(\psi_t) = X^{\mathrm{main}}$;

$t_{\mathrm{RGAend}} := t$;

Therefore, we can find an embedding ψ' extending ψ_{RGA} which embeds X_i^{buf} to $V_i \setminus \mathrm{Im}(\psi_{\mathrm{RGA}})$. Repeating this for each $i \in [r]$, which we may do since X^{buf} is an independent set in H, we obtain the desired embedding of H into G. $\qquad\square$

5.2 Proof of the degenerate graph RGA lemma

We prove Lemma 5.2 by analysing Algorithm 4. The analysis is quite similar to what we saw before in the proofs of Lemmas 3.1 and 4.1 (indeed, the main difference is that we are more careful to bound powers of p using D rather than just Δ), so we will be brief and highlight the differences. The only difference between Algorithm 3 and Algorithm 4 is that vertices of X^e are embedded into V^{c} rather than V^{main} or V^{q}.

    We will see that (*ORD* 1') (see Definition 5.1 of (D,p,m)-bounded order) is what we need to make Lemma 2.27 work with properties $\mathrm{NS}(\varepsilon, r_1, D)$ and $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$, so allowing us to prove that

$B_t(x)$ is always much smaller than $A_t(x)$. As in the proof of Lemma 4.1, our first task is to show that the algorithm a.a.s. completes successfully. Again by Lemma 2.29 we can show that the queue remains small. We can also show, using (*ORD* 2'), that all vertices $x \in X_i \setminus X^e$ which enter the queue have $\pi^\tau(x) \le D/2$. We embed the vertices of $X^e \cap X_i$ which enter the queue greedily into V_i^c, and there are so few such vertices that this is guaranteed to succeed. Property $\mathrm{LCON}(\varepsilon, r_1, D)$ turns out to be what we need to verify that the queue embedding of the remaining vertices a.a.s. is successful, using the same strategy as in the proof of Lemma 4.1.

At this stage we have (*DRGA* 1) simply because the algorithm completes, while (*DRGA* 2) follows from Lemma 2.29. It remains to prove (*DRGA* 3), which is where we need to use (*ORD* 3'). Here we deviate from the strategy we saw previously. We can no longer assume that neighbours of buffer vertices appear early in τ, and thus we have to prove that for any $v \in V(G)$, even towards the end of the embedding, it is still reasonably likely that neighbours of buffer vertices are embedded to $N_G(v)$. We will see (Claim 5.7 below) that properties $\mathrm{NS}(\varepsilon, r_1, D)$, $\mathrm{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$ and $\mathrm{LCON}(\varepsilon, r_1, D)$ allow us to show that a.a.s. given any $\ell \le D - \max_{x \in X^{\mathrm{main}} \setminus X^e} \pi^\tau(x)$ vertices $v_1, \dots, v_\ell$ of G such that $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$ is not small, the set $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$ is never completely filled by $\mathrm{Im}(\psi_t)$. The idea of the proof remains similar to that of Lemma 2.30. We show, as there, that when the first small fraction of the vertices in the order τ are embedded, only at most half of $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$ is covered. But then we repeat this, showing that the next small fraction of τ covers only at most half of what remains, and so on, so that when all vertices are embedded what remains is an exponentially small, but bounded away from zero, fraction of the original $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$. Once we have shown this, completing the proof that $N_H(x)$ is not too unlikely to be embedded to $N_G(v)$ for any $x \in X^{\mathrm{buf}}$ and $v \in V(x)$ can be done along similar lines to the proof of Lemma 2.31.

Proof of Lemma 5.2. We require

$$\mu \le \tfrac{1}{8}, \, \rho \le \frac{\mu^2 \zeta^2 d^{2D}}{10000(\Delta+1)}, \, \varepsilon' \le \frac{\mu \zeta d^{\Delta+1}}{1000 \Delta^2 \kappa 4^\Delta} 2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}, \, \varepsilon \le \frac{d\varepsilon'}{\kappa}$$

$$\text{and} \quad p \ge \frac{10000 \cdot 4^{10\Delta^3} \Delta^3 \kappa r_1^2}{\varepsilon \mu \zeta d^{\Delta+1}} 2^{280(D+1)\mu^{-1}\zeta^{-1}d^{-D}} \left(\frac{\log n}{n}\right)^{1/D}.$$

We run Algorithm 4, using the order τ supplied to Lemma 5.2. We claim that we can apply Lemma 2.27 to bound $B_{\tau(x)-1}(x)$. To see this, we verify properties (*i*)–(*iv*) of Lemma 2.27. We consider three cases.

If x has no neighbours after $\tau(x)$ in τ, then Lemma 2.27 states that $B_{\tau(x)-1}(x) = \emptyset$ without requiring any property of Γ.

If x has a neighbour y with $\tau(x) < \tau(y)$, then (*ORD* 1') states that $\pi^\tau(x) \le D-1$, so in particular $\pi^*_{\tau(x)-1}(x) \le D-1$, satisfying (*i*). Furthermore, (*ORD* 1') states that $\pi^\tau(y) \le D$, so $\pi^*_{\tau(x)-1}(y) \le D-1$ since x is not embedded, satisfying (*ii*). Now suppose that $xy, yz \in E(H)$ for some unembedded y and z. There are two cases to consider. First, if $\tau(y) < \tau(z)$ then (*ORD* 1') states that $\pi^\tau(y) \le D-1$ and $\pi^\tau(z) \le D$. Since x and y are unembedded at time $\tau(x)-1$, we thus have $\pi^*_{\tau(x)-1}(y) \le D-2$ and $\pi^*_{\tau(x)-1}(z) \le D-1$, as required by (*iii*). Second, if $\tau(y) > \tau(z)$, then $\pi^*_{\tau(x)-1}(z) \le \pi^\tau(z) \le D-1$ by (*ORD* 1'), while since $\pi^\tau(y) \le D$ and both the vertices x and z are unembedded, we have $\pi^*_{\tau(x)-1}(y) \le D-2$, again as required by (*iii*).

Lastly, if there are unembedded vertices y and z such that xy, yz, xz are all edges of H, then we claim $\pi^*_{\tau(x)-1}(x), \pi^*_{\tau(x)-1}(y), \pi^*_{\tau(x)-1}(z) \leq D-2$. To see this, suppose $\tau(y) < \tau(z)$. By (*ORD* 1') we have $\pi^\tau(x) \leq D-2$, $\pi^\tau(y) \leq D-1$ and $\pi^\tau(z) \leq D$. Since x and y are unembedded, we have $\pi^*_{\tau(x)-1}(y) \leq D-1-1 = D-2$, and $\pi^*_{\tau(x)-1}(z) \leq D-2$. Thus condition (*iv*) is satisfied.

We conclude that we can apply Lemma 2.27, and hence obtain $\left|B_{\tau(x)-1}(x)\right| \leq 20\Delta^2 \varepsilon' p^{\pi^\tau(x)} |V(x)|$ for each $x \in X^{\mathrm{main}}$. Next, we show that certain invariants are maintained during the running of Algorithm 4. These are identical to the invariants seen in the previous two RGA lemmas, but are repeated here for the reader's convenience.

Claim 5.3. *The following hold at each time t in the running of Algorithm 4.*

(*INV* 1) *ψ_t is a good partial embedding,*

(*INV* 2) *Either $|A^{\mathrm{main}}_t(x)| \geq \frac{1}{2}\mu(d-\varepsilon')^{\pi^*_t(x)} p^{\pi^*_t(x)} |V^{\mathrm{main}}(x)|$ or $x \in Q_t$.*

(*INV* 3) *When we embed x to create ψ_{t+1}, we do so uniformly at random into a set of size at least $\frac{1}{10}\mu\zeta(dp)^{\pi^*_t(x)}|V(x)|$.*

Proof. Algorithm 4 maintains (*INV* 1) and (*INV* 2) by definition. Since

$$\tfrac{1}{2}\mu(d-\varepsilon')^{\pi^\tau(x)} p^{\pi^\tau(x)} - 20\Delta^2 \varepsilon' p^{\pi^\tau(x)} > \tfrac{1}{10}\mu\zeta d^{\pi^\tau(x)} p^{\pi^\tau(x)}$$

by choice of ε' for each $x \in X^{\mathrm{main}}$, by Lemma 2.27 it also maintains (*INV* 3). □

As in the previous proofs, the conditions of Lemma 2.29 are met, so we conclude that with probability at least $1 - 2^{-n/(\kappa r_1)} r_1$, for each $i \in [r]$ and time t when Algorithm 4 is running, we have $\left|Q_t \cap X_i\right| < \rho|X_i| + |X^*_i| \leq 2\rho|X_i|$ by (*G* 6).

We now want to show that Algorithm 4 a.a.s. does not halt with failure. Since we know that $\left|B_{\tau(x)-1}(x)\right| < 20\Delta^2 \varepsilon' p^{\pi^\tau(x)} |V(x)|$, by (*GPE* 3) and (*G* 5) it is enough to show that a.a.s. we have

$$\left|A^{\mathrm{q}}_{\tau(x)-1}(x)\right| \geq \tfrac{1}{2}\left|C^{\mathrm{q}}_{\tau(x)-1}(x)\right| \quad \text{for} \quad x \in X^{\mathrm{main}} \cap Q_{\tau(x)-1} \setminus X^e, \quad \text{and} \qquad (5.1)$$

$$\left|A^{\mathrm{c}}_{\tau(x)-1}(x)\right| \geq \tfrac{1}{2}\left|C^{\mathrm{c}}_{\tau(x)-1}(x)\right| \quad \text{for} \quad x \in X^{\mathrm{main}} \cap Q_{\tau(x)-1} \cap X^e.$$

The latter of these is easy, since we embed only vertices of X^e into V^{c}, and $|X^e| \leq \frac{1}{2}\varepsilon p^{\max_{x \in X^e} \pi^\tau(x)} n/r_1$, which by (*GPE* 3) and choice of ε is smaller than $\frac{1}{2}\left|C^{\mathrm{c}}_{\tau(x)-1}(x)\right|$. In the next three claims we show (5.1) holds.

First, we show that we need only concern ourselves with x such that $\pi^\tau(x) \leq D/2$.

Claim 5.4. *For each $x \in X^{\mathrm{main}} \setminus X^e$, if $\pi^\tau(x) > D/2$ then there is no time t such that $x \in Q_t$.*

Proof. We require $\varepsilon' \leq \frac{1}{4}d^D\mu$ and $\varepsilon \leq \kappa^{-1}\varepsilon'$.

Suppose $\pi^\tau(x) > D/2$ and $x \notin X^e$. Let i be such that $x \in X_i$. By (*ORD* 2'), x is not image restricted, and all $\pi^\tau(x)$ neighbours of x which precede x in τ occur in the $\frac{1}{2}\varepsilon p^{\pi^\tau(x)} n/r_1 \leq \varepsilon' p^{\pi^\tau(x)} |X_i|$ places of τ before x. Because x is not image restricted, at time $t \leq \tau(x) - \varepsilon' p^{\pi^\tau(x)} |X_i|$ we have $C^{\mathrm{main}}_t(x) = V^{\mathrm{main}}_i$ for some i, and hence $A^{\mathrm{main}}_t(x) = V^{\mathrm{main}}_i \setminus \mathrm{Im}(\psi_t)$. Because $|X^{\mathrm{main}}_i| = |V^{\mathrm{main}}_i| - \mu|V_i|$, it follows that $|A^{\mathrm{main}}_t(x)| \geq \mu|V_i|$, and in particular x does not enter Q_t at time t.

Now suppose $t = \tau(x) - \varepsilon' p^{\pi^{\tau}(x)}|X_i| + s$ for some $0 \le s \le \varepsilon' p^{\pi^{\tau}(x)}|X_i| - 1$. We claim that

$$|A_t^{\mathrm{main}}(x)| \ge (dp - \varepsilon' p)^{\pi_t^*(x)} \mu |V_i| - s.$$

Indeed, for $s = 0$ we have $\pi_t^*(x) = 0$ and the statement $|A_t^{\mathrm{main}}(x) \ge \mu|V_i|$ was established above. Now suppose the statement holds for some $s \ge 0$. At time $t = \tau(x) - \varepsilon' p^{\pi^{\tau}(x)}|X_i| + s$ we embed a vertex y to obtain ψ_{t+1}. If y is neither a neighbour of x nor in X_i, then $A_{t+1}^{\mathrm{main}}(x) = A_t^{\mathrm{main}}(x)$ and the statement holds for $s+1$. If y is a neighbour of x, then we embed y to a vertex of $A_t(y) \setminus B_t(y)$, and in particular to a vertex v such that $\deg_G\left(v; A_t^{\mathrm{main}}(x)\right) \ge (d - \varepsilon') p |A_t^{\mathrm{main}}(x)|$. We conclude that

$$|A_{t+1}^{\mathrm{main}}(x)| \ge (d - \varepsilon') p\left((dp - \varepsilon' p)^{\pi_t^*(x)} \mu |V_i| - s\right) > (dp - \varepsilon' p)^{\pi_{t+1}^*(x)} \mu |V_i| - s - 1.$$

as desired. Finally, if $y \in X_i$ then we have $A_{t+1}^{\mathrm{main}}(x) = A_t^{\mathrm{main}}(x) \setminus \{y\}$, and again

$$|A_{t+1}^{\mathrm{main}}(x)| \ge (dp - \varepsilon' p)^{\pi_{t+1}^*(x)} \mu |V_i| - s - 1.$$

Thus for each $\tau(x) - \varepsilon' p^{\pi^{\tau}(x)}|X_i| \le t \le \tau(x) - 1$ we have

$$|A_t^{\mathrm{main}}(x)| \ge (dp - \varepsilon' p)^{\pi_t^*(x)} \mu |V_i| - \varepsilon' p^{\pi^{\tau}(x)}|X_i| \ge \tfrac{1}{2}(dp - \varepsilon' p)^{\pi_t^*(x)} \mu |V_i|,$$

where the final inequality is by choice of ε'. It follows that x does not satisfy the condition to enter Q_t. Since we embed x at time $\tau(x) - 1$, we conclude that as desired there is no time t such that $x \in Q_t$. $\qquad\square$

As in the proof of Lemma 4.1 the next step is to establish a 'sum condition'. In fact, we establish almost the same inequality as (4.6) in Claim 4.5 (the right hand side is identical, while the left hand side sum runs over all vertices outside X^e, including those y with $J_y \ne \emptyset$), though we use a different analysis to do so.

Claim 5.5. *Suppose that Γ has* $\mathrm{LCON}(\varepsilon, r_1, D)$. *Then the following statement holds.*
For any t, any $1 \le \ell \le D/2$, any $i \in [r]$ and any vertices $v_1, \dots, v_\ell$ such that

$$\deg_\Gamma(v_1, \dots, v_\ell; V_i) \ge \varepsilon' p^\ell |V_i|$$

we have

$$\sum_{\substack{y \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e: \\ \tau(y) \le t}} \frac{\left|U_{\tau(y)-1}(y) \cap N_\Gamma^*(v_1, \dots, v_\ell; V_i)\right|}{\left|A_{\tau(y)-1}^{\mathrm{q}}(y) \setminus B_{\tau(y)-1}(y)\right|} \le \frac{\mu \zeta d^\Delta \deg_\Gamma(v_1, \dots, v_\ell; V_i)}{20}. \tag{5.2}$$

The proof of this amounts to checking that the images under our embedding of the $N_H(y)$ for $y \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e$ form a family of sets to which we can apply the LCON property. This property gives us the desired inequality.

Proof. We require $\rho \le \frac{\mu^2 \zeta^2 d^{D+\Delta}}{10000(\Delta+1)}$, $\varepsilon' \le \mu \zeta (d/4)^D$ and $\varepsilon \le \frac{\varepsilon'}{\kappa}$. Since the left hand side of (5.2) is increasing in t, we may assume t is the time at which Algorithm 4 ends.

Given $1 \le \ell \le D/2$ and $i \in [r]$, let $v_1, \dots, v_\ell$ be vertices such that $U := N_\Gamma^*(v_1, \dots, v_\ell; V_i)$ has size at least $\varepsilon' p^\ell |V_i|$. By choice of ε and ε', we have $|U| \ge \varepsilon p^\ell n / r_1$. It follows that we can apply $\mathrm{LCON}(\varepsilon, r_1, D)$

to bound the number of edges in $\mathrm{CG}(\Gamma,U,\mathcal{F})$ for any $1 \leq j \leq D/2$ and family $\mathcal{F}$ of pairwise disjoint j-sets in $V(\Gamma) \setminus U$. We set

$$\mathcal{F}_j := \left\{ \Pi_{\tau(x)-1}(x) \cup J_x : x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e, \pi^\tau(x) = j, \tau(x) \leq t \right\}.$$

Observe that $\mathcal{F}_j$ is indeed a family of j-sets by definition of $\pi^\tau(x)$, that these j-sets are pairwise disjoint because no vertex of H has more than one neighbour in X_i by $(H2)$, and element F of $\mathcal{F}_j$ intersects $U \subseteq V_i$ because these sets F are images under ψ of neighbours of vertices in X_i. Because $\varepsilon|U|,|X_i^{\mathrm{main}} \cap Q_t| \leq 2\rho|X_i|$, by choice of ε and by $\mathrm{LCON}(\varepsilon,r_1,D)$ we have

$$e\big(\mathrm{CG}(\Gamma,U,\mathcal{F}_j)\big) \leq 7p^j|U|(2\rho|X_i|) = 14\rho p^j|U||X_i|.$$

By definition of $U_t(x)$ (see Section 2.3.5) we have the following equality

$$e\big(\mathrm{CG}(\Gamma,U,\mathcal{F}_j)\big) = \sum_{x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e: \pi^\tau(x)=j, \tau(x) \leq t} |U_{\tau(x)-1}(x) \cap U|$$

and hence

$$\sum_{x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e: \pi^\tau(x)=j, \tau(x) \leq t} |U_{\tau(x)-1}(x) \cap U| \leq 14\rho p^j|U||X_i|$$

for each $1 \leq j \leq D/2$. Observe that the $j = 0$ case of the same inequality is trivially true (indeed with 14 replaced by 2, though we will not use this): each vertex of $X_i^{\mathrm{main}} \cap Q_t$ contributes at most $|U|$ to the sum. Now by $(INV3)$ we have

$$\left| A_{\tau(x)}^{\mathrm{q}}(x) \setminus B_{\tau(x)}(x) \right| \geq \tfrac{1}{10}\mu\zeta(dp)^{\pi^\tau(x)}|V_i|$$

for each $x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e$ with $\tau(x) \leq t$. Finally, since any vertex x in $X_i^{\mathrm{main}} \cap Q_t \setminus X^e$ satisfies $\pi^\tau(x) \leq D/2$ by Claim 5.4, we conclude that

$$\sum_{\substack{x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e: \\ \tau(x) \leq t}} \frac{|U_{\tau(x)}^{\mathrm{q}}(x) \cap U|}{|A_{\tau(x)}^{\mathrm{q}}(x) \setminus B_{\tau(x)}(x)|} \leq \sum_{j=0}^{D/2} \sum_{\substack{x \in X_i^{\mathrm{main}} \cap Q_t \setminus X^e: \\ \pi^\tau(x)=j, \tau(x) \leq t}} \frac{|U_{\tau(x)}^{\mathrm{q}}(x) \cap U|}{\tfrac{1}{10}\mu\zeta(dp)^j|V_i|}$$

$$\leq \sum_{j=0}^{D/2} \frac{14\rho p^j|U||X_i|}{\tfrac{1}{10}\mu\zeta(dp)^j|V_i|} \leq \frac{14(1+D/2)\rho|U||X_i|}{\tfrac{1}{10}\mu\zeta d^{D/2}|V_i|} \leq \frac{\mu\zeta d^\Delta}{20}|U|,$$

where the final inequality is by choice of ρ, as desired. □

We can now complete the proof that candidate sets in V^{q} are not substantially covered by $\mathrm{Im}(\psi_t)$ for any t.

Claim 5.6. *A.a.s. for each $x \in X^{\mathrm{main}} \setminus X^e$, at each time $t \leq \tau(x) - 1$ and before the termination of Algorithm 4, we have $\left| C_t^{\mathrm{q}}(x) \cap \mathrm{Im}(\psi_t) \right| < \tfrac{1}{2}\left| C_t^{\mathrm{q}}(x) \right|$.*

Proof. The proof of Claim 4.6, if we replace X^{main} with $X^{\mathrm{main}} \setminus X^e$, and replace (4.6) with (5.2), without handling image restricted vertices specially (since condition (5.2) deals with them), gives this claim verbatim. □

By the same logic as in the proof of Lemma 4.1, if the good event of Claim 5.6 holds, it follows that candidate sets in V^q are not significantly covered by $\mathrm{Im}(\psi_t)$, i.e. the inequality (5.1) holds. This implies that Algorithm 4 a.a.s. completes successfully. Again, the successful running implies (*DRGA* 1), and the good event of Lemma 2.29 holding implies (*DRGA* 2).

We would like to emphasise at this point that we have established all the desired conclusions of Lemma 5.2 except (*DRGA* 3), and we have so far not used condition (*ORD* 3'), and not used the second part of (*ORD* 1') (which states $\pi^\tau(x) \le D_x - 1$ for $x \in N(\tilde{X}) \setminus X^{\mathrm{buf}}$). It follows that one can establish a version of Lemma 5.2 omitting (*DRGA* 3) given an order τ which need only satisfy the first part of (*ORD* 1') (which states $\pi^\tau(x) \le D_x$ for all $x \in V(H)$) and (*ORD* 2'). We will return to this point in Section 7.2.3.

We now turn to proving (*DRGA* 3) in the next two claims. First, in Claim 5.7, we show that G-common neighbourhoods of at most $D - \max_{x \in X^{\mathrm{main}} \setminus X^e} \pi^\tau(x)$ vertices which are large do not get completely filled at any time in the running of Algorithm 4. Note that the upper bound on ℓ in Claim 5.7 comes from (*ORD* 3'). What the good event of Claim 5.7, together with (*ORD* 3'), says is the following. If $v = v_1 \in V_i$ is any vertex (which we eventually want to show is candidate for many vertices of X_i^{buf}) and $v_2, \dots, v_\ell$ are the embedded neighbours of some $x \in X_i^{\mathrm{buf}}$, then the G-common neighbourhood of $v_1, \dots, v_\ell$ is not covered by $\mathrm{Im}(\psi_t)$ provided $t \le \tau(x) - \frac{1}{2}\varepsilon p^D n / r_1$. We will see that neighbours embedded after this time can be dealt with easily. Claim 5.7 has two parts, (*a*) dealing with candidate sets in V^{main} and (*b*) dealing with candidate sets in V^q. The reason for this is that neighbours of a buffer vertex x may end up in the queue, and we need to show that even in this case v is reasonably likely to be a candidate for x. This claim replaces Lemma 2.30, which we used in the proofs of the previous two RGA lemmas.

Claim 5.7. *Suppose that Γ has $\mathrm{LCON}(\varepsilon, r_1, D)$. Then a.a.s. the following holds at each time t. Given any $i \in [r]$, any $1 \le \ell \le D - \max_{x \in X^{\mathrm{main}} \setminus X^e} \pi^\tau(x)$, and any vertices $v_1, \dots, v_\ell$ of G, the following hold.*

(*a*) *If $\deg_G(v_1, \dots, v_\ell; V_i^{\mathrm{main}}) \ge (dp/2)^\ell |V_i^{\mathrm{main}}|$, then*

$$\left| N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}}) \setminus \mathrm{Im}(\psi_t) \right| \ge$$
$$(1 - 4\mu)2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}} \deg_G(v_1, \dots, v_\ell; V_i^{\mathrm{main}}) . \tag{5.3}$$

(*b*) *If $\deg_G(v_1, \dots, v_\ell; V_i^q) \ge (dp/2)^\ell |V_i^q|$, then*

$$\left| N_G^*(v_1, \dots, v_\ell; V_i^q) \setminus \mathrm{Im}(\psi_t) \right| \ge \tfrac{1}{2} \deg_G(v_1, \dots, v_\ell; V_i^q) . \tag{5.4}$$

The proof of (5.3), as discussed at the beginning of this section, roughly amounts to repeating the argument used to prove Lemma 2.30 several times. We show that, for some small constant η, each successive interval consisting of an η-fraction of the vertices X_i^{main} is likely to cover less than half of whatever of $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$ was uncovered before embedding that interval. The argument for each given interval is similar to that seen in Lemma 2.30, though here we take a short-cut by using property LCON rather than NS, which simplifies the calculations.

Proof. Let $h = \max_{x \in X^{\mathrm{main}} \setminus X^e} \pi^\tau(x)$. We require

$$\mu \le \tfrac{1}{8}, \qquad\qquad \rho \le \frac{\mu^2 \zeta d^D}{1000(D+1)},$$

$$\varepsilon < \tfrac{1}{2}2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}\left(\tfrac{d}{2}\right)^{D+1} \quad \text{and} \qquad p \ge \tfrac{1000\Delta^2}{\varepsilon\mu\zeta d^D}\left(\tfrac{\log n}{n}\right)^{1/D} .$$

We first prove inequality (5.3). Set $\eta = \frac{\mu \zeta d^D}{280(D+1)}$. Given $i \in [r]$ and $1 \le \ell \le D - h$, suppose that $v_1, \dots, v_\ell \in V(\Gamma)$ are such that $N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}}) \ge (dp/2)^\ell |V_i^{\mathrm{main}}|$. We split the vertices of $X_i^{\mathrm{main}} \setminus X^e$ into intervals $\mathrm{Int}_1, \dots, \mathrm{Int}_{1/\eta}$ of equal size, with the first being the first $\eta |X_i^{\mathrm{main}} \setminus X^e|$ vertices in the order τ, and so on.

We now aim to show the following.

(‡) For any fixed set $U \subseteq V_i$ of size at least $\varepsilon p^\ell |V_i|$ and any $1 \le j \le 1/\eta$, conditioning on the embedding up to the last vertex of Int_{j-1}, with probability at least $1 - n^{-D-1}$ there are at most $\frac{1}{2}|U|$ vertices of Int_j are embedded to U.

To that end, for each $x \in \mathrm{Int}_j$ let $\mathscr{H}_{x,j}$ denote the history up to but not including the embedding of x. By definition of h, all vertices $x \in \mathrm{Int}_j$ have $\pi^\tau(x) \le h$, so we split the vertices x of Int_j up into $h + 1$ classes according to $\pi^\tau(x)$. We apply $\mathrm{LCON}(\varepsilon, r_1, D)$ with the set U and the family $\mathcal{F}_s$ of s-sets given by the embedded images of $N_H(x)$ for $x \in \mathrm{Int}_j$ with $\pi^\tau(x) = s$. Since $\varepsilon |U|, |\mathcal{F}_s| \le \eta |X_i^{\mathrm{main}}|$, we obtain the inequality

$$\sum_{x \in \mathrm{Int}_j : \pi^\tau(x) = s} \left| U_{\tau(x)-1}(x) \cap U \right| \le 7 p^s |U| \eta |X_i^{\mathrm{main}}|.$$

Now each $x \in \mathrm{Int}_j$ with $\pi^\tau(x) = s$ is by ($INV\,3$) embedded uniformly at random into a set of size at least $\frac{1}{10} \mu \zeta (dp)^s |V_i|$. We thus have

$$\sum_{x \in \mathrm{Int}_j} \mathbb{P}\left(x \text{ is embedded to } U \big| \mathscr{H}_{x,j}\right) \le (h+1) \frac{70 \eta |U|}{\mu \zeta d^h} \le (D+1) \frac{70 \eta |U|}{\mu \zeta d^D}$$

and by Lemma 2.2 with $\delta = 1$ we see that the probability that more than

$$140(D+1) \eta \mu^{-1} \zeta^{-1} d^{-D} |U| \le \tfrac{1}{2}|U|$$

of the vertices in Int_j are embedded to U is at most

$$\exp\left(-(D+1) \frac{70 \eta |U|}{3 \mu \zeta d^D}\right) \le n^{-D-1},$$

where the inequality is because $|U| \ge \varepsilon p^{D-h} |V_i|$ and by choice of p. This completes the proof of (‡).

Let $U_0 = N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$, and $U_j = U_{j-1} \setminus \psi_{t_{j+1}}(\mathrm{Int}_j)$ for each $j \ge 1$. In the event that $|U_j| \ge \frac{1}{2}|U_{j-1}|$ for each j, we have

$$|U_j| \ge 2^{-1/\eta} (d/2)^\ell p^\ell |V_i^{\mathrm{main}}| \ge \varepsilon p^\ell |V_i|$$

for each j, where the final inequality is by choice of ε. In other words, each U_{j-1} is large enough to play the rôle of U in (‡), so taking a union bound over the $1/\eta$ choices of j, we see that with probability at least $1 - \eta^{-1} n^{-D-1}$ we have

$$\left| N_G^*(v_1, \dots, v_\ell; V_i^{\mathrm{main}}) \setminus \mathrm{Im}(\psi_{t_j}) \right| \ge (1 - 4\mu) 2^{1-j} \deg_G(v_1, \dots, v_\ell; V_i^{\mathrm{main}})$$

for each $1 \le j \le 1/\eta + 1$. Taking a union bound over the at most r_1 choices of $i \in [r]$ and $n + n^2 + \dots + n^D$ choices of $v_1, \dots, v_\ell$ we see that the desired event of the lemma holds with probability at least $1 - r_1 D \eta^{-1} n^{-1}$, which tends to one as n tends to infinity. This completes the proof of (5.3).

The proof of (5.4) is simpler. Observe that $|X_i^q| \leq 2\rho|X_i| < \eta|X_i^{\mathrm{main}}|$, and $\deg_G(v_1,\ldots,v_\ell;V_i^q) \geq (dp/2)^\ell|V_i^q| \geq \varepsilon p^\ell|V_i|$ so that the analysis above, with only one 'interval' $X_i^q \setminus X^e$, gives the desired result.

□

Finally, we are in a position to prove (*DRGA* 3), which we do in the following claim. The proof is quite similar to that of Lemma 2.31, which amounts to showing that any vertex $v \in V_i$ is likely to be a candidate vertex for 'reasonably' many of the buffer vertices X_i^{buf}. More precisely, our strategy is to fix a vertex $v \in V_i$ and show that the inequality (5.5) below, which encapsulates these reasonable bounds, holds with sufficiently high probability to apply the union bound over all $v \in V(G)$. To do this, we introduce the concept of an $(x \to v)$-buffer-friendly partial embedding for $x \in X_i^{\mathrm{buf}}$. As in the proof of Lemma 2.31, this is a good partial embedding which has a few extra properties that let us show that it is not too unlikely that $N_H(x)$ is embedded to $N_G(v)$. As there, the proof that it is not too unlikely that at each step Algorithm 4 maintains an $(x \to v)$-buffer-friendly partial embedding is mainly 'bookkeeping' and long but not very hard. However, there is an important difference. It is no longer useful to simply multiply the conditional probabilities of maintaining an $(x \to v)$-buffer-friendly partial embedding in order to estimate the conditional probability that $N_H(x)$ is embedded to $N_G(v)$. This is because the vertices $N_H(x)$ are no longer embedded as a segment of τ, so that Lemma 2.2 is not applicable directly to these products of conditional probabilities. Nevertheless we show below that we can still apply Lemma 2.2, several times and with some extra care, to obtain the desired bounds.

Claim 5.8. *The following holds a.a.s. for each $i \in [r]$. Let b be such that X_i^{buf} is a degree-b buffer. Then for each $v \in V_i$ we have*

$$\left|\{x \in X_i^{\mathrm{buf}} : v \in C_{t_{\mathrm{RGAend}}}(x)\}\right| \geq 4^{-10\Delta^3} 2^{-1000D^2\mu^{-1}\zeta^{-1}d^{-D}} d^{2D^2} \Delta^{-3} \mu p^b |X_i|. \tag{5.5}$$

Proof. We require

$$\mu \leq \tfrac{1}{8}, \qquad \varepsilon' < \frac{\mu\zeta d^\Delta}{1000\Delta^2\kappa 4^\Delta} \cdot 2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}, \qquad \varepsilon \leq \varepsilon',$$

and
$$p > 10000 \cdot 4^{10\Delta^3}\Delta^3\kappa r_1\mu^{-1}d^{-\Delta-1}\cdot 2^{280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}\left(\tfrac{\log n}{n}\right)^{1/D}.$$

Let $\beta = (1-4\mu)2^{-280(D+1)\mu^{-1}\zeta^{-1}d^{-D}}$. It is convenient to consider only vertices of X_i^{buf} which are far from X^e; since X^e is very small, doing so does not exclude many vertices of X_i^{buf}.

Given $x \in X_i^{\mathrm{buf}}$ which is at distance greater than three from any vertex of X^e, and $v \in V_i$, we say that a good partial embedding ψ is an $(x \to v)$-*buffer-friendly partial embedding* if the following hold.

(*BPE* 1) For each $y \in \mathrm{Dom}(\psi) \cap N_H(x)$ we have $\psi(y) \in N_G(v)$.

(*BPE* 2) For each unembedded $y \in N_H(x)$ we have

$$|U(y) \cap N_\Gamma(v)| = (p \pm \varepsilon' p)^{\pi^*(y)} \deg_\Gamma(v;V(y)),$$
$$|U^{\mathrm{main}}(y) \cap N_\Gamma(v)| = (p \pm \varepsilon' p)^{\pi^*(y)} \deg_\Gamma(v;V^{\mathrm{main}}(y)) \quad \text{and}$$
$$|U^q(y) \cap N_\Gamma(v)| = (p \pm \varepsilon' p)^{\pi^*(y)} \deg_\Gamma(v;V^q(y))$$

(*BPE* 3) For each unembedded $y \in N_H(x)$ we have

$$|C^q(y) \cap N_G(v)| \geq (dp - \varepsilon' p)^{\pi^*(y)} \deg_G(v;V^q(y)) \quad \text{and}$$
$$|C^{\mathrm{main}}(y) \cap N_G(v)| \geq (dp - \varepsilon' p)^{\pi^*(y)} \deg_G(v;V^{\mathrm{main}}(y))$$

($BPE\,4$) For each unembedded $y, z \in N_H(x)$ with $yz \in E(H)$, the pair $\big(U(y) \cap N_\Gamma(v), U(z) \cap N_\Gamma(v)\big)$ is $(\varepsilon_{\pi^*(y),\pi^*(z)}, d, p)$-regular in G.

($BPE\,5$) For each unembedded $y \in N_H(x)$ and $z \in N_H(y)$, we have that the pair $\big(U(y) \cap N_\Gamma(v), U(z)\big)$ is $(\varepsilon_{\pi^*(y),\pi^*(z)}, d, p)$-regular in G.

Much as in the proof of Lemma 2.31, the empty partial embedding ψ_0 is an $(x \to v)$-buffer-friendly partial embedding. Indeed, ($BPE\,1$)–($BPE\,3$) are trivially true since by ($H\,5$) neighbours of buffer vertices are not image restricted. For ($BPE\,4$) and ($BPE\,5$), since buffer vertices are by ($H\,5$) at distance at least three from image restricted vertices the sets $U(y)$ and $U(z)$ are equal to $V(y)$ and $V(z)$ respectively, and the required regularity is thus given by ($G\,2$), since $X_i^{\mathrm{buf}} \subseteq \tilde{X}_i$.

Our aim now is to give lower bounds on the probability that, given that at some time t we currently have an $(x \to v)$-buffer-friendly partial embedding, Algorithm 4 embeds the next vertex in a way that maintains an $(x \to v)$-buffer-friendly partial embedding. To that end, given an $(x \to v)$-buffer-friendly partial embedding ψ and an unembedded vertex y, we let $P(y)$ be the set of *poor vertices for y*, namely those vertices $u \in C(y)$ such that either $\psi \cup \{y \to u\}$ is not an $(x \to v)$-buffer-friendly partial embedding, or such that there is an unembedded $z \in N_H(y) \cap N_H(x)$ such that either of the following conditions hold:

$$
\begin{aligned}
\deg_G\big(u; A^{\mathrm{main}}(z) \cap N_G(v)\big) &\leq (d - \varepsilon')p \big| A^{\mathrm{main}}(z) \cap N_G(v)\big| \quad \text{or} \\
\deg_G\big(u; A^{\mathrm{q}}(z) \cap N_G(v)\big) &\leq (d - \varepsilon')p \big| A^{\mathrm{q}}(z) \cap N_G(v)\big|.
\end{aligned}
\tag{5.6}
$$

Observe that if y is at distance four or greater from x in H, then $P(y)$ is always empty. This has two important consequences. First, we are about to talk about the probability of no vertices being embedded to poor vertices: this is really a condition on at most $3\Delta^3$ vertex embeddings. Second, it means that if $x, x' \in X_i^{\mathrm{buf}}$, by ($H\,2$) the distance between x and x' is at least ten, so that any given vertex embedding affects at most one of having an $(x \to v)$-buffer-friendly partial embedding and having an $(x' \to v)$-buffer-friendly partial embedding.

Before we try to estimate the probability of maintaining an $(x \to v)$-buffer-friendly partial embedding, we strengthen the conclusion of Claim 5.7 to cover the 'last few vertices' before embedding $y \in N_H(x)$.

Fact 5.9. *For a given $x \in X_i^{\mathrm{buf}}$ at distance greater than three from X^e and a fixed $v \in V(x)$, provided that up to time t no vertex has been embedded to a poor vertex (with respect to an $(x \to v)$-buffer-friendly partial embedding), for each $y \in N_H(x)$ with $\tau(y) > t$, the following a.a.s. hold.*

$$
\begin{aligned}
\big| A_t^{\mathrm{main}}(y) \cap N_G(v)\big| &\geq \tfrac{\beta}{2}(dp/2)^{\pi_t^*(y)} \deg_G\big(v; V^{\mathrm{main}}(y)\big) \quad \text{and} \\
\big| A_t^{\mathrm{q}}(y) \cap N_G(v)\big| &\geq \tfrac{1}{4}(dp/2)^{\pi_t^*(y)} \deg_G\big(v; V^{\mathrm{q}}(y)\big).
\end{aligned}
\tag{5.7}
$$

We will see that Claim 5.7, together with ($ORD\,3'$), show that this fact holds provided $t \leq \tau(y) - \frac{1}{2}\varepsilon p^D n/r_1$. The number of vertices embedded in the remaining $\frac{1}{2}\varepsilon p^D n/r_1$ steps is too small to significantly fill up either of these sets, so we need only show that embedding neighbours of y does not adversely affect (5.7). This is guaranteed by condition (5.6) in the definition of poor vertices.

Proof. Let $h = \max_{z \in X^{\mathrm{main}} \setminus X^e} \pi^\tau(z)$. Since $X_i^{\mathrm{buf}} \subseteq \tilde{X}$ and $y \in N_H(x)$ we infer by ($ORD\,3'$) that all but at most s neighbours z of y satisfy $\tau(y) - \tau(z) \leq p^D \cdot \frac{1}{2}\varepsilon n/r_1$ for some $s \leq D - 1 - h$ (recall that τ is a $(D, p, \frac{1}{2}\varepsilon n/r_1)$-bounded order). Therefore, if $t \leq \tau(y) - \frac{1}{2}\varepsilon p^D n/r_1$ then $\pi_t^*(y) \leq s \leq D - 1 - h$,

so that $\Pi_t^*(y) \cup \{v\}$ is by *(BPE 3)*, *(G 3)* and *(G 1)* a set of $s+1 \leq D-h$ vertices in G with at least $(dp/2)^{s+1}|V^q(y)|$ common neighbours in $V^q(y)$ and at least $(dp/2)^{s+1}|V^{\mathrm{main}}(y)|$ common neighbours in $V^{\mathrm{main}}(y)$.

Therefore, by Claim 5.7, we have

$$\left|A_t^{\mathrm{main}}(y) \cap N_G(v)\right| \geq \beta (dp/2)^{\pi_t^*(y)} \deg_G\left(v; V^{\mathrm{main}}(y)\right) \quad \text{and}$$
$$\left|A_t^q(y) \cap N_G(v)\right| \geq \tfrac{1}{2}(dp/2)^{\pi_t^*(y)} \deg_G\left(v; V^q(y)\right),$$

which gives (5.7) for this range of t. Now suppose that

$$t = \tau(y) - \varepsilon p^D \cdot \tfrac{1}{2} n/r_1 + \ell$$

for some $0 \leq \ell < \tfrac{1}{2}\varepsilon p^D n/r_1$. We claim that

$$\left|A_t^{\mathrm{main}}(y) \cap N_G(v)\right| \geq \beta (dp/2)^{\pi_t^*(y)} \deg_G\left(v; V^{\mathrm{main}}(y)\right) - \ell \quad \text{and}$$
$$\left|A_t^q(y) \cap N_G(v)\right| \geq \tfrac{1}{2}(dp/2)^{\pi_t^*(y)} \deg_G\left(v; V^q(y)\right) - \ell. \tag{5.8}$$

Indeed, we have just established that (5.8) holds for $\ell = 0$. Given $\ell \geq 1$, suppose that (5.8) holds for $\ell - 1$, and consider the vertex z embedded at time $\tau(z) = \tau(y) - \varepsilon p^D \cdot \tfrac{1}{2} n/r_1 + \ell$.

If z is not a neighbour of y, then its embedding decreases the sizes of the sets on the left hand side of (5.8) by at most one, so the inequality continues to hold. If z is a neighbour of y, then by (5.6) in the definition of $P(z)$ and since, by the assumption of Fact 5.9 no vertex has been embedded to a poor vertex, we have

$$\left|A_{\tau(z)}^q(y) \cap N_G(v)\right| \geq (d - \varepsilon')p\left|A_{\tau(z)-1}^q(y) \cap N_G(v)\right|$$
$$\geq \tfrac{1}{2}(dp/2)^{\pi_{\tau(z)}^*(y)} \deg_G\left(v; V^q(y)\right) - (d - \varepsilon')p\ell$$

where the second inequality is by choice of ε' and (5.8). A similar inequality holds for $\left|A_{\tau(z)}^{\mathrm{main}}(y) \cap N_G(v)\right|$. Thus again (5.8) continues to hold. We conclude that (5.8) holds for all $0 \leq \ell < \tfrac{1}{2}\varepsilon p^D n/r_1$.

Now since $x \in X^{\mathrm{buf}}$, we have $\tau(x) > \tau(y)$, so that by *(ORD 1')* we have $\pi_t^*(y) \leq \pi^\tau(y) \leq D-1$. By *(G 3)* and choice of ε, we conclude that (5.7) holds for all $t < \tau(y)$ as desired. □

We now continue with the proof of Claim 5.8 by estimating the probability of not choosing a poor vertex at a given step in Algorithm 4. Recall that $B(y)$ is the set of bad vertices for a given good partial embedding and unembedded y, as defined in Definition 2.26.

Fact 5.10. *Suppose that ψ is a good partial embedding generated together with a queue Q by Algorithm 4 which is also an $(x \to v)$-buffer-friendly partial embedding, where x is at distance greater than three from any vertex of X^e, and (5.7) holds for ψ and all unembedded y. Given an unembedded y, suppose that u is a vertex chosen uniformly at random in $A^{\mathrm{main}}(y) \setminus B(y)$, or in $A^q(y) \setminus B(y)$, or in $A^c(y) \setminus B(y)$, according to whether $y \notin Q \cup X^e$, or $y \in Q \setminus X^e$, or $y \in X^e$ respectively. Then the following hold.*

(a) *If $\mathrm{dist}_H(x,y) > 3$ then $\mathbb{P}\big(u \notin P(y)\big) = 1$.*

(b) *If $\mathrm{dist}_H(x,y) \in \{2,3\}$ then $\mathbb{P}\big(u \notin P(y)\big) \geq \tfrac{1}{2}$.*

(c) If $\text{dist}_H(x,y) = 1$ then $\mathbb{P}(u \notin P(y)) \geq \beta d^{D-1} 4^{-D} p$.

The proof of this fact is quite similar to the analysis in Lemma 2.31. We remark that the use we make of ψ and Q being generated by Algorithm 4 is that the invariants of Claim 5.3 hold; we do not in the following proof perform any analysis of the probabilistic process generating ψ and Q.

Proof. If $\text{dist}_H(x,y) > 3$ then $P(y) = \emptyset$, so (a) is trivial. From now on we assume $\text{dist}_H(x,y) \leq 3$.

By ($INV\,3$) we embed y uniformly at random into a set whose size is at least $\frac{1}{10}\mu\zeta(dp)^{\pi^\tau(y)}|V(y)|$. Thus in order to show (b) we simply need to establish that $P(y) \cup B(y)$ is small compared to this set; this can be established using $\text{NS}(\varepsilon, r_1, D)$, $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$ and the buffer and good partial embedding properties, as in the proof of Lemma 2.31. We need to use the fact, given by ($ORD\,1$'), that $\pi^\tau(z) \leq D_z - 1$ for each $z \in N(\tilde{X}) \setminus X^{\text{buf}}$, in other words that $\pi^\tau(z) \leq D - 2$ for all $z \in N_H(x)$ and $\pi^\tau(z) \leq D - 3$ if there is a triangle zww' in H with $\tau(w), \tau(w') > \tau(z)$. In order to establish that few vertices fail (5.6), we need to know that $A^{\text{main}}(z) \cap N_G(v)$ and $A^{\text{q}}(z) \cap N_G(v)$ are not small in comparison to $U(z) \cap N_\Gamma(v)$. This follows from (5.7) since $z \in N_H(x)$. Since the proof is otherwise as in the proof of Lemma 2.31, we omit further details.

Finally we come to (c). We embed y uniformly at random into a subset of $U(y)$, which by ($GPE\,2$) has size at most $2p^{\pi^\tau(y)}|V(y)|$. If $xy \in H$, we need to show that it is not too unlikely that y is embedded to $N_G(v)$. By (5.7), ($G\,3$) and choice of β, we see that

$$\left|A^{\text{main}}(y) \cap N_G(v)\right|, \left|A^{\text{q}}(y) \cap N_G(v)\right| \geq \tfrac{\beta}{4}(dp/2)^{\pi^\tau(y)} \deg_G(v; V(y)).$$

It is thus enough to show that at most half of these vertices are in $P(y) \cup B(y)$, which again we can do using $\text{NS}(\varepsilon, r_1, D)$, $\text{RI}(\varepsilon, (\varepsilon_{a,b}), \varepsilon', d, r_1, D)$ and the buffer and good partial embedding properties, as in the proof of Lemma 2.31. Again, we use the fact $\pi^\tau(z) \leq D_z - 1$ for $z \in N_H(x)$ in this verification. We omit the details. □

We now continue with the proof of Claim 5.8. We fix i and $v \in V_i$. Our goal is to show that with very high probability a reasonable fraction of the $x \in X_i^{\text{buf}}$ have $N_H(x)$ embedded to $N_G(v)$, since this immediately implies the claim. We have established in Fact 5.10 that for any given vertex x of X_i^{buf} far from X^e, the probability that $N_H(x)$ is embedded to $N_G(v)$ is not too small. Fact 5.10 states that this probability is more or less given by the embeddings of vertices at distance at most three from x, and distinct vertices of X_i^{buf} are at distance at least ten by ($H\,2$), so that no vertex is within distance three of two distinct $x, x' \in X_i^{\text{buf}}$. In the proof of Lemma 2.31 at this point we simply apply Lemma 2.2 to complete the proof. However, here we cannot do this, because the sets $N_H(x)$ interleave each other in the order τ. We have to be a bit more careful, as follows.

For each $x \in X_i^{\text{buf}}$, let M_x be the set of vertices of H at distance one, two or three from x. We say *x survives at step j* if after the embedding of the jth vertex of M_x in the order τ, we still have an $(x \to v)$-buffer-friendly partial embedding. We will use Lemma 2.2 to show that a reasonable fraction of $x \in X_i^{\text{buf}}$ survive at step 1. We would like then to use Lemma 2.2 again to show that (because a reasonable fraction of vertices survive at step 1) a reasonable fraction of vertices survive at step 2, and so on.

In order to carry this out, it is convenient not in fact to look at all $x \in X_i^{\text{buf}}$, but only at a subset Y in which the probability of surviving at each step j (as given by Fact 5.10) does not depend on the particular vertex in $x \in Y$ but only on j. We now construct such a set Y. For any $x \in X_i^{\text{buf}}$ at distance greater than

three from X^e, there are at most $\Delta^3 + \Delta^2 + \Delta < 3\Delta^3$ vertices of H at distance at most 3 from x in H. We can therefore associate to each $x \in X_i^{\text{buf}}$ at distance greater than three from X^e a 0–1 vector with at most $3\Delta^3$ non-zero entries, taking the value 0 at place j if the jth vertex y of M_x (in the order τ) is not a neighbour of x, and 1 if y is a neighbour of x. There are at most $3\Delta^3 2^{3\Delta^3}$ choices for this vector; we fix a most common choice $\mathbf{c} = (c_i)$, and let Y be the subset of vertices of X_i^{buf} at distance at least three from X^e associated to this most common choice. By construction, by $(H\,3)$, by choice of ε and since $|X^e| \leq \frac{1}{2}\varepsilon p^D n/r_1$ and $|X_i| \geq n/\kappa r_1$, we have

$$|Y| \geq \tfrac{1}{6}\Delta^{-3}2^{-3\Delta^3}|X_i^{\text{buf}}|. \tag{5.9}$$

For each $x \in Y$ and $1 \leq j \leq |\mathbf{c}|$ (where $|\mathbf{c}|$ is the number of 1-entries of $\mathbf{c}$), let $\mathscr{H}_{x,j}$ be the history of Algorithm 4 up to the point immediately before embedding the jth vertex of M_x. For each $x \in Y$, we create a collection of Bernoulli random variables $A_{x,1}$ for $x \in Y$, which are set equal to one if either x survives at step 1 or (5.7) has failed before embedding the first element of M_x, and zero otherwise. We identify in each M_x the first (in τ) element $w_{x,1}$, and sort the $A_{x,1}$ according to the order induced by τ on the $w_{x,1}$. In this order, the $A_{x,1}$ are sequentially dependent. We define

$$s_1 = \begin{cases} \frac{1}{2}\beta d^{D+1}4^{-D}p|Y| & \text{if} \quad c_1 = 1 \\ \frac{1}{4}|Y| & \text{if} \quad c_1 = 0 \end{cases}.$$

Now for each $A_{x,1}$, either Fact 5.10 gives us a lower bound on $\mathbb{E}[A_{x,1}|\mathscr{H}_{x,1}]$, or (5.7) has failed before we embed the first vertex of M_x, in which case $\mathbb{E}[A_{x,1}|\mathscr{H}_{x,1}] = 1$. We conclude, by definition of s_1, that

$$\sum_{x \in Y} \mathbb{E}[A_{x,1}|\mathscr{H}_{x,1}] \geq 2s_1.$$

By Lemma 2.2, with $\delta = 1/4$, with probability at least $1 - e^{-s_1/24}$ we have $\sum_{x \in Y} A_{x,1} \geq s_1 + 1$. Observe that if (5.7) does not fail for any y or t (which is a.a.s. the case by Fact 5.9) then $\mathcal{A}_1$ holding says that at least $s_1 + 1$ vertices survive at step 1.

There remains a small technical difficulty in continuing. We would like to count vertices surviving at step 2, and use Lemma 2.2 and a lower bound on the sum of conditional expectations given by the product of s_1 and the probability bound from Fact 5.10 to show that the number of vertices surviving at step 2 is likely to be large. But it is possible that less than s_1 vertices survive at step 1, so that this bound on the sum of conditional expectations does not hold almost surely (it only holds a.a.s.).

To get around this problem, we use the following trick. Given x, immediately after embedding the first vertex of M_x in Algorithm 4, let g be the number of vertices x' in Y such that $A_{x,1}$ is certainly equal to one, i.e. such that the first vertex of $M_{x'}$ was embedded to a vertex which is neither poor nor bad, so x' survives at step 1. Let h be the number of vertices x' in Y such that the first vertex of $M_{x'}$ has not yet been embedded. If $g + h \leq s_1$, we say x is *dangerous at step* 1 if the following holds. In other words, we say x is dangerous at step 1 if after embedding the first vertex of M_x we already know that $\mathcal{A}_1$ does not occur.

We now define Bernoulli random variables $A_{x,2}$ for $x \in Y$, set equal to one if either x survives at step 2, or x is dangerous at step 1, or (5.7) has failed for some y prior to embedding the second vertex of M_x, and zero otherwise. The point of this definition is that it gives us an *a priori* lower bound on the sum of conditional expectations of the $A_{x,2}$, as we require to apply Lemma 2.2, but nevertheless with very high

probability the sum of the $A_{x,2}$ does simply count the number of vertices in Y which survive at step 2, because with very high probability (5.7) does not fail and $\mathcal{A}_1$ does occur, so no vertex is dangerous at step 1. Again, this collection of random variables has a natural order given by the order τ on the second vertices of the M_x. Observe that we know whether x is dangerous at step 1 before we embed the second vertex of M_x. Thus in this order the random variables are sequentially dependent. This justifies that we can apply Lemma 2.2 to estimate the sum of the $A_{x,2}$.

We now say what exactly this application of Lemma 2.2 gives us, and explain how to analyse steps $3, \ldots, |\mathbf{c}|$ in the same way. We need to give integers s_j for $2 \le j \le |\mathbf{c}|$, which are our desired lower bounds on the number of vertices in Y surviving at step j, to define the events $\mathcal{A}_j$ and the concept of *dangerous at step j* for $j \ge 2$, and define the random variables $A_{x,j}$ for $3 \le j \le |\mathbf{c}|$.

For each $2 \le j \le |\mathbf{c}|$, we set

$$s_j = \begin{cases} \frac{1}{2}\beta d^{D+1}4^{-D}ps_{j-1} & \text{if} \quad c_j = 1 \\ \frac{1}{4}s_{j-1} & \text{if} \quad c_j = 0 \end{cases}.$$

We let the event $\sum_{x \in Y} A_{x,j} \ge s_j + 1$ be $\mathcal{A}_j$. We say $x \in Y$ is dangerous at step j if immediately after embedding the jth vertex of M_x, the number of $x' \in Y$ such that $A_{x',j}$ is certainly equal to one, plus the number of $x' \in Y$ such that the jth vertex of $M_{x'}$ has not yet been embedded, is at most s_j. Again, this means that after embedding the jth vertex of M_x, we already know $\mathcal{A}_j$ does not occur. Finally, for each $j \ge 3$ we define the Bernoulli random variables $A_{x,j}$ for $x \in Y$, set equal to one if x survives at step j, or is dangerous at step $j-1$, or (5.7) has failed before embedding the jth vertex of M_x, and zero otherwise.

Before we continue, we observe that

$$s_1 \ge s_2 \ge \cdots \ge s_{|\mathbf{c}|} \ge \left(\frac{1}{4}\right)^{\Delta^2 + \Delta^3} \left(\frac{1}{2}\beta d^{D+1}4^{-D}p\right)^b |Y|, \tag{5.10}$$

where b is the degree of the buffer X_i^{buf}, i.e. $\deg_H(x) = b$ for each $x \in X_i^{\text{buf}}$; this number exists by $(H\,6)$. The reason is simply that there are at most $\Delta^2 + \Delta^3$ vertices of H at distance two or three in H from any given $x \in Y$, and b neighbours of x. In particular, by choice of p, by our lower bound on $|Y|$ and since n is sufficiently large, we have $3s_j/2 > s_j + 1$ for each $1 \le j \le |\mathbf{c}|$.

We now show that with very high probability $\mathcal{A}_j$ occurs for each $1 \le j \le |\mathbf{c}|$. Recall that by Fact 5.10, we have the lower bound $\sum_{x \in Y} \mathbb{E}[A_{x,1}|\mathcal{H}_{x,1}] \ge 2s_1$. By Lemma 2.2 with $\delta = 1/4$, and the observation $3s_1/2 > s_1 + 1$, the probability that $\mathcal{A}_1$ fails is thus at most $e^{-s_1/24}$.

Now for $j \ge 2$, for each $x \in Y$, either x survived at step $j-1$ and (5.7) has not failed before embedding the jth vertex of M_x (in which case the expectation of $A_{x,j}$ conditioned on $\mathcal{H}_{x,j}$ is at least the quantity given in Fact 5.10), or x was dangerous at step $j-1$, or (5.7) failed before embedding the jth vertex of M_x (in which case we have $\mathbb{E}[A_{x,h}|\mathcal{H}_{x,j}] = 1$), or none of these occur (in which case we have $\mathbb{E}[A_{x,h}|\mathcal{H}_{x,j}] \ge 0$). Furthermore, by definition of the $A_{x,j}$ and of the event $\mathcal{A}_{j-1}$, we know (a priori, before beginning the embedding) that one of the first three cases occurs for at least s_{j-1} of the vertices Y, so that we have $\sum_{x \in Y} \mathbb{E}[A_{x,j}|\mathcal{H}_{x,j}] \ge 2s_j$ by definition of s_j. Applying Lemma 2.2, with $\delta = 1/4$, and since $3s_j/2 > s_j + 1$, we conclude that the probability that $\mathcal{A}_j$ fails is at most $e^{-s_j/24}$.

By definition, if (5.7) never fails and each event $\mathcal{A}_1, \ldots, \mathcal{A}_{|\mathbf{c}|}$ occurs, then no vertex is dangerous at any step. Thus the number of vertices of Y surviving at step $|\mathbf{c}|$ is at least $s_{|\mathbf{c}|}$. By definition, if x survives

at step $|\mathbf{c}|$ then $N_H(x)$ is embedded to $N_G(v)$, so that v is a candidate for x. Since $s_{|\mathbf{c}|}$ is by (5.10), (5.9) and $(H\,3)$ at least

$$\left(\tfrac{1}{4}\right)^{\Delta^2+\Delta^3}\left(\tfrac{1}{2}\beta d^{D+1}4^{-D}p\right)^b \cdot \tfrac{1}{6}\Delta^{-3}2^{-3\Delta^3} \cdot 4\mu|X_i|$$

we conclude that (5.5) holds for v, as desired. The probability that any given one of the events $\mathcal{A}_1,\ldots,\mathcal{A}_{|\mathbf{c}|}$ fails is at most $e^{-s_{|\mathbf{c}|}/24} < n^{-D-1}$ by (5.10), (5.9), $(H\,3)$, since $|X_i| \geq n/\kappa r_1$ and by choice of p. Taking the union bound over the at most $3\Delta^3$ events $\mathcal{A}_j$, and the at most n choices of v, and since a.a.s. (5.7) never fails by Fact 5.9, we conclude that a.a.s. (5.5) holds for all $v \in V_i$ for each $i \in [r]$, completing the proof of Claim 5.8. $\qquad\square$

The good event of Claim 5.8 holding gives $(DRGA\,3)$. Since a.a.s. the good events of each of the above claims and lemmas hold, this completes the proof of Lemma 5.2. $\qquad\square$

Chapter 6

Proofs of applications

In this chapter we prove the various theorems listed in Section 1.2. We indicate when stronger results are proved (subsequently) elsewhere. To keep this section brief we often only sketch proofs.

6.1   Universal graphs

We begin by showing universality of $G(n, p)$ for (almost) spanning graphs with bounded degree and bounded degeneracy.

Proof of Theorem 1.1. We first show (a). Observe that any d-degenerate n-vertex graph has at most dn edges and so contains at least $\frac{n}{2d+1}$ vertices of degree at most $2d$. We will apply Lemma 1.23 with input Δ, with $\Delta_{R'} = 8\Delta$, $D = 2d + 1$, $\alpha = \frac{\gamma}{10d}$, $\zeta = 1$, density $\frac{1}{2}$, and $\kappa = 2$. Lemma 1.23 returns $\varepsilon > 0$, and for input $r_1 = 8\Delta$ also C. Choose $\varepsilon^* \ll \varepsilon$ and suppose C is also large enough for Lemma 2.17(a) with input ε^*, 8Δ and 2. Let $p \geq C\left(\frac{\log n}{n}\right)^{1/(2d+1)}$. Fix an equipartition $V_1, \ldots, V_{8\Delta}$ of $[n]$. Now $G(n, p)$ a.a.s. satisfies the good event of Lemma 1.23. By Lemma 2.17(a) it a.a.s. has property $NS(\varepsilon^*, 8\Delta, D)$. Finally, using Theorem 2.1 and the union bound it is easy to check that a.a.s. for each i and $v \in V_i$, the vertex v has $(1 \pm \varepsilon)p|V_j|$ neighbours in V_j for each $j \neq i$. Fix a graph Γ with each of these properties. Property $NS(\varepsilon^*, 8\Delta, D)$ implies that any pair of disjoint subsets of $V(\Gamma)$, each of size at least $\frac{pn}{16\Delta}$, is $(\varepsilon, \frac{1}{2}, p)$-regular, so letting $R = R' = K_{8\Delta}$ we have an $(\varepsilon, \frac{1}{2}, p)$-regular R-partition which has super-regularity and one- and two-sided inheritance on R'.

   Given H on n vertices with degeneracy d and maximum degree Δ, let $X_1, \ldots, X_{8\Delta}$ be an equipartition of $V(H)$ into independent sets, with each X_i containing at least $\frac{1}{8\Delta(2d+1)}|X_i|$ vertices of degree at most $2d$. This equipartition exists by Lemma 2.4. We designate the vertices of degree at most $2d$ as potential buffer vertices. We do not image restrict any vertices. Then (perhaps after reordering) the X_i form an R-partition of $V(H)$ which is size-compatible with the V_i. We let τ be a degeneracy order on $V(H)$. We move all buffer vertices to the end of the ordering τ. Observe that this slightly changed ordering satisfies conditions $(ORD\,1')$–$(ORD\,3')$ and thus is $(D, p, \varepsilon^* n/r_1)$-bounded. Then the conditions of Lemma 1.23 are satisfied, so $H \subseteq \Gamma$ as desired.

   To obtain (b), we work identically except that we add to each part of $V(H)$ isolated vertices to obtain a size-compatible R-partition, and we designate these isolated vertices as potential buffer vertices instead

of the low-degree vertices in H. Then Lemma 1.23, applied with $D = 2d$, $p = C(\frac{\log n}{n})^{1/D}$ and all other parameters staying as before, gives the desired conclusion. $\qquad\square$

It is quite easy to use Lemma 1.25 to show that sufficiently pseudorandom graphs are universal for bounded-degree graphs.

Sketch proof of Theorem 1.2. Given G, we take a random equipartition of $V(G)$ into $\Delta + 1$ clusters $V_1, \ldots, V_{\Delta+1}$. For each $v \in V(G)$ and i, the quantity $\deg(v; V_i)$ is hypergeometrically distributed with mean at least $\frac{1}{4}p|V_i|$, so by Theorem 2.1 and the union bound we see that with positive probability we have $\deg(v; V_i) \geq \frac{1}{8}p|V_i|$ for each v and i. We fix such a partition, and let $R = R' = K_{\Delta+1}$. Using the Hajnal-Szemerédi theorem (Theorem 2.3) we can for any n-vertex graph H with $\Delta(H) \leq \Delta$ find a size-compatible R-partition of $V(H)$, and applying Lemma 1.25 with all vertices designated as potential buffer vertices, with $d = \frac{1}{16}$ and with no vertices image restricted, we see that $H \subseteq G$ as desired. $\qquad\square$

6.2   Partition universality

In this section we prove Theorems 1.3, 1.4 and 1.5. We need Ramsey's Theorem and Turán's Theorem. The Ramsey number $r_k(K_t)$ is the smallest integer n such that no matter how one colours the edges of K_n with k colours, there is a monochromatic copy of K_t. Ramsey [70] proved these numbers exist, while the following quantitative statement is due to Erdős and Szekeres [37].

Theorem 6.1 (Erdős and Szekeres [37])**.** *For any k and t we have $r_k(K_t) \leq k^{kt}$.*

Turán, generalising a result of Mantel [67], proved the following.

Theorem 6.2 (Turán [79])**.** *For any $r \geq 3$, any n-vertex K_r-free graph has at most $\frac{(r-2)n^2}{2(r-1)} + o(n^2)$ edges.*

Finally we state the version of the sparse regularity lemma for many colours that we are going to apply. We also say that a graph G with density p is (η, D)-*upper-uniform* if for all disjoint sets U and W of cardinality at least $\eta v(G)$ we have $e_G(U, W) \leq Dp|U||W|$.

Lemma 6.3 (Sparse regularity lemma, coloured version [50])**.** *For any real $D, \varepsilon > 0$, integers k and t_0, there exist $\eta > 0$ and T such that any graph G of edge density p and on at least t vertices, which is (η, D)-upper-uniform and whose edges are coloured with k colours admits a partition of $V(G)$ into V_1, $\ldots$, V_t with the following properties.*
(a) $t_0 \leq t \leq T$*.*
(b) $\big||V_i| - |V_j|\big| \leq 1$ *for all i, j.*
(c) *all but at most εt^2 pairs (V_i, V_j) are (ε, d, p)-fully-regular in each of the k colours for some possibly different d.*

The assumption on (η, D)-upper uniformity is easily seen to be satisfied for any moderate η and $D > 1$ by our random graph $G(n, p)$ and by the bijumbled graphs that we will be considering. The partition asserted by Lemma 6.3 is called ε-regular. We define a coloured multigraph $R(\delta)$ on $[t]$ (associating each $i \in [t]$ with the class V_i) as follows. We put an edge ij into $E(R)$ in colour c if the pair (V_i, V_j) is (ε, δ, p)-regular in colour c in G.

We now prove that $G(n,p)$ is a.a.s. r-partition universal for $\mathcal{H}(n,d,\Delta)$ provided $p \geq C\left(\frac{\log n}{n}\right)^{1/(2d)}$, which improves on the result of Kohayakawa, Rödl, Schacht and Szemerédi [54] for graphs with $d \leq \Delta/2$. A proof of the result of [54] can be obtained along very similar lines.

Proof of Theorem 1.3. We apply Lemma 1.23 with input Δ, with $\Delta_{R'} = 0$, $D = 2d$, $\alpha = \frac{1}{2}$, $\zeta = 1$, density $\frac{1}{2d}$, and $\kappa = 1$. Lemma 1.23 returns $\varepsilon > 0$, which we suppose is small enough for the application of Turán's Theorem below. We let r_1 be large enough for the k-coloured sparse regularity lemma with input ε and also for the applications of Turán's and Ramsey's Theorem below, and obtain $C \geq r_1$ from Lemma 1.23. Given p, we generate $\Gamma = G(Cn, p)$, suppose it satisfies the good event of Lemma 1.23 and take any k-colouring of its edges. We apply the sparse regularity lemma for coloured graphs, Lemma 6.3, to this coloured graph. We obtain a coloured reduced graph with at most an ε-fraction of pairs not forming edges, in which we find a $4^{\Delta+1}$-vertex clique by Turán's Theorem, and in that a monochromatic $(\Delta+1)$-vertex clique by Ramsey's Theorem. Thus we have $\Delta + 1$ equal-sized clusters $V_1, \dots, V_{\Delta+1}$, each of size at least n, such that there is a colour c in which each pair of clusters is $(\varepsilon, \frac{1}{2r}, p)$-regular. Let G be the subgraph of $\Gamma[V_1, \dots, V_{\Delta+1}]$ containing all edges of colour c, let $R = K_{\Delta+1}$ and R' be the empty graph.

Given $H \in \mathcal{H}(n, \Delta)$, by Theorem 2.3 we can find an equipartition of $V(H)$ into $\Delta + 1$ independent sets, each of size at most $\lceil n/(\Delta+1) \rceil < n/2$. We 'pad' each set by adding at least $n/2$ isolated vertices to obtain a size-compatible R-partition $X_1, \dots, X_{\Delta+1}$, and designate the isolated vertices as potential buffer vertices which come last in the degeneracy order τ of H. Then the result follows by Lemma 1.23, with τ being a $(D, p, \varepsilon n/r)$-bounded order on $V(H)$. $\qquad\square$

Sketch proof of Theorem 1.4. The same approach as for Theorem 1.3 works, replacing Lemma 1.23 with Lemma 1.25. It is easy to check that the upper-uniformity condition of Lemma 6.3 is satisfied for Γ with the bijumbledness condition of the theorem. $\qquad\square$

Sketch proof of Theorem 1.5. We follow the same approach as in the proof of Theorem 1.3, with the exceptions that we use $p = C\left(\frac{\log n}{n}\right)^{1/\Delta}$, that we use Lemma 1.21 instead of Lemma 1.23, and that after generating $\Gamma = G(n, p)$ which satisfies the good event of Lemma 1.21, we form G by deleting a minimum number of edges to remove all copies of $K_{2\Delta, 2\Delta}$. The expected number of copies of $K_{2\Delta, 2\Delta}$ in $G(n, p)$ is $O\left(p^{4\Delta^2} n^{4\Delta}\right) = o(pn^2)$, so by Markov's inequality we see that a.a.s. we delete $o(pn^2)$ edges of Γ to form G. This G is the claimed r-partition universal graph for $\mathcal{H}(n, \Delta)$. The remainder of the proof of Theorem 1.3 proves the partition universality, since the number of edges deleted is too small to destroy regularity of any pair of clusters. $\qquad\square$

6.3 Maker-Breaker games

The proof of Theorem 1.6 uses the result of Ferber, Krivelevich and Naves [38], which informally says that if $p \ll b^{-1}$ is not too small, there is a (randomised) Maker strategy for the $1 : b$ game on K_n which gives Maker a subgraph of $G(n, p)$ with minimum degree very close to pn. More formally, they prove the following.

Theorem 6.4. *For any $\frac{\log n}{n} \ll p \ll b^{-1}$, there is a randomised Maker strategy which, for any fixed strategy of Breaker in the $1:b$ game on K_n, generates a graph Γ from the distribution $G(n,p)$, and a spanning subgraph G of edges obtained by Maker which a.a.s. has minimum degree $(1-o(1))pn$.*

This reduces the proof to showing a (far from optimal in terms of edge deletions) local resilience result for the graph classes we consider.

Sketch proof of Theorem 1.6. We begin by sketching the proof of (b), i.e. that Maker obtains $\mathcal{H}'(n,\Delta)$-universality. We take an equipartition of $[n]$ into $\Delta+1$ clusters $V_1, \ldots, V_{\Delta+1}$. We let $\varepsilon > 0$ be sufficiently small for Lemma 1.21, with no vertices image restricted, and we set $\varepsilon^* \ll \varepsilon$. Suppose C is large enough for Lemma 1.21, and $p = C\left(\frac{\log n}{n}\right)^{1/\Delta}$. Let $c > 0$ be sufficiently small, and suppose $b \leq c\left(\frac{n}{\log n}\right)^{1/\Delta}$. Then a.a.s. in the $1:b$ game on K_n, Maker obtains a subgraph G of $\Gamma = G(n,p)$ with $\delta(G) \geq (1-\varepsilon^*)pn$ by Theorem 6.4. It is easy to check, using Chernoff's inequality, that each vertex has $(1 \pm 2\varepsilon^*)pn/(\Delta+1)$ neighbours in both Γ and G in each set V_i.

Since the total number of edges of Γ which are not in G leaving any vertex is at most $4\varepsilon^* pn$, the total number of such edges between any two vertex sets X and Y is at most $4\varepsilon^* pn|X|$. An easy application of the Chernoff bound shows that (X, V_i) is $\left(\varepsilon^*, \frac{3}{4}, p\right)$-regular in Γ for each V_i and disjoint vertex set X with $|X| \geq pn/(4\Delta)$, and since $\varepsilon^* \ll \varepsilon$ we conclude that any such (X, V_i) is also $\left(\varepsilon, \frac{1}{2}, p\right)$-regular in G.

Letting $R = R' = K_{\Delta+1}$, we see that G is $\left(\varepsilon, \frac{1}{2}, p\right)$-super-regular on R' and has one-sided inheritance on R'. Given any graph $H \in \mathcal{H}'(n,\Delta)$ we use Theorem 2.3 to find an equipartition of $V(H)$ into $\Delta+1$ independent sets $X_1, \ldots, X_{\Delta+1}$. Then the conditions of Lemma 1.21 are satisfied, so we have $H \subseteq G$ as desired.

For (a), i.e. the almost-spanning $\mathcal{H}(n,\Delta)$-universality game, we repeat the same argument, replacing $[n]$ with $[(1+\delta)n]$ and 'padding' each equipartition class of H with independent vertices to be size-compatible with $\mathcal{V}$. We take these independent vertices to be the potential buffer vertices, so that two-sided inheritance is not needed.

For (c) and (d), the degeneracy statements, we replace Lemma 1.21 with Lemma 1.23, taking respectively $D = 2d$ for the almost-spanning universality and $D = 2d+1$ for the spanning universality. Again, similar to Theorem 1.3, it is easy to verify that the degeneracy order of $H \in \mathcal{H}(n,d,\Delta)$ is an appropriately bounded order in the sense of Definition 1.22.

In both cases, since we have shown that Maker has a randomised strategy which wins with positive probability against any strategy of Breaker, it follows that Breaker does not have a winning strategy. Since finite draw-free games are determined, we conclude that Maker does have a winning strategy. $\quad\square$

Note that, by an analysis similar to the proof of Theorem 1.3, one can show that Maker also succeeds with this strategy in making a graph which is *r-partition universal* for $\mathcal{H}(cn,\Delta)$ (respectively, for $\mathcal{H}(cn,d,\Delta)$) for some small $c > 0$, matching the density of the best known constructions.

6.4   Resilience for low-bandwidth graphs

For this section we need the minimum degree form of the sparse regularity lemma, which we quote from [25], the paper in which the bipartite case of Theorem 1.9 is proved. To state it we need to define two concepts. First, an ε-*equipartition* of a vertex set V is a partition $V = V_0 \dot\cup \cdots \dot\cup V_r$ such that $|V_0| \leq \varepsilon|V|$

and $|V_1| = \cdots = |V_r|$. Second, if G is a graph with vertex set V, then the (ε, d, p)-*lower-regular reduced graph of* G, with respect to a given ε-equipartition $V = V_0 \dot{\cup} \cdots \dot{\cup} V_r$, is the graph on $[r]$ with edges ij corresponding to (ε, d, p)-lower-regular pairs (V_i, V_j) in G.

Lemma 6.5 ([25], Lemma 4.4). *For all $\vartheta \in [0, 1]$, $\varepsilon > 0$ and every integer r_0, there exists $r_1 \geq 1$ such that for all $d \in [0, 1]$ the following holds a.a.s. for $\Gamma = G(n, p)$ if $\log^4 n / (pn) = o(1)$. Let $G = (V, E)$ be a spanning subgraph of Γ with $\deg_G(v) \geq \vartheta \deg_\Gamma(v)$ for all $v \in V$. Then there is an ε-equipartition of G with (ε, d, p)-lower-regular reduced graph R of minimum degree $\delta(R) \geq (\vartheta - d - \varepsilon)v(R)$, and $r_0 \leq v(R) \leq r_1$.*

Using this, we can sketch the proof of Theorem 1.9. The strategy consists of modifying the argument in [28], the paper in which the Bandwidth Theorem was proved. We will not state formally the lemmas from that paper which we require; the reader not familiar with that argument will wish to read the following sketch in conjunction with Section 2 of [28], in which the lemmas are formally stated and their use outlined. The changes to their strategy we make are as follows. Their 'Lemma for G' is replaced with Lemma 6.5 (since we do not need most of the properties of the 'Lemma for G' in this setting), and Theorem 1.8 finds a 'backbone graph' in the resulting (ε, d, p)-reduced graph R. We can use their 'Lemma for H' as written, and in this setting it gives a partition of $V(H)$ which is directly suitable to apply the blow-up lemma, Lemma 1.21, yielding the desired embedding of H into G.

Sketch proof of Theorem 1.9. Given $\gamma > 0$ and Δ, we choose $d \ll \gamma$ and $r_0 \gg \Delta$. We apply Lemma 1.21 with input Δ, $\Delta_{R'} = \Delta_J = 0$, $\alpha = \gamma/2$, $\zeta = 1$, d and $\kappa = 2$. Lemma 1.21 returns $\rho, \varepsilon > 0$, of which we are only interested in ε. We assume, without loss of generality, that $\varepsilon \ll d$. We let r_1 be returned by Lemma 6.5 for input $\vartheta = \frac{\chi(H) - 1}{\chi(H)} + \gamma$, ε and r_0. Finally, we choose $\beta \ll r_1^{-1}$, and let C be returned by Lemma 1.21 for input r_1.

Now, given $p \geq C \left(\frac{\log n}{n} \right)^{1/\Delta}$, we generate $\Gamma = G(n, p)$, and assume it satisfies the conditions of Lemmas 1.21 and 6.5 for the parameters given above. Let a graph H on $(1 - \gamma)n$ vertices with $\Delta(H) \leq \Delta$ and $\mathrm{bw}(H) \leq \beta n$ be given. Let $r = \chi(H) \leq \Delta + 1$. Let a spanning subgraph G of Γ with minimum degree $\vartheta pn = \left(\frac{r-1}{r} + \gamma \right)pn$ be given.

We apply Lemma 6.5 to G, obtaining an (ε, d, p)-reduced graph R, with $r_0 \leq v(R) \leq r_1$, and minimum degree at least $\left(\frac{r-1}{r} + \frac{\gamma}{2} \right)v(R)$. By Theorem 1.8, we can find in R a spanning *backbone graph*: This consists of a collection of vertex-disjoint copies of K_r, which come in a linear order, such that between one copy of K_r and the next there is a copy of $K_{r,r}$ with a perfect matching removed. This graph is r-colourable and has maximum degree $3(r - 1) \leq 3\Delta$ and bandwidth at most $2r \leq 2\Delta$, so that Theorem 1.8 is indeed applicable provided r_0 is sufficiently large compared to Δ.

It is quite easy to find a homomorphism from H to the backbone graph. We simply divide $V(H)$ up into intervals in the bandwidth order, and map successive intervals of $V(H)$ to successive copies of K_r in the backbone graph, choosing vertices of each copy of K_r according to a fixed r-colouring of H. The rôle of the bandwidth restriction here is to ensure that edges of H either lie within one interval, or go from one interval to the next, so that we only need edges in R from one copy of K_r in the backbone graph to the next in order to obtain a homomorphism. The point of fixing an r-colouring of H is that the ith vertex in one copy of K_r and that in the next are not adjacent in the backbone graph (since a perfect matching was removed from the $K_{r,r}$ between them) and we need to ensure that no edge of H will be assigned to have one endpoint in each.

Unfortunately, this is not quite enough: the colour classes of H could be quite unbalanced, so that the homomorphism we have just described maps many more vertices of H to some vertices of R than others. In order to repair this, we need to 'rebalance', which requires that each copy of K_r in the backbone graph extends to a copy of K_{r+1} using some other vertex of R (which may be anywhere in the backbone graph). The Lemma for H of [28] now states that given a backbone graph whose r-cliques extend to copies of K_{r+1}, and H, there is a homomorphism f from $V(H)$ to $V(R)$ in which each vertex of R is the image of approximately the same number of vertices of H. Since $\varepsilon \ll \gamma$, since $v(H) \leq (1 - \gamma)n$ and by choice of β sufficiently small, the 'approximately' in this statement in particular guarantees $\left| f^{-1}(i) \right| \leq \left(1 - \frac{\gamma}{2}\right) |V_i|$ for each $i \in V(R)$.

It remains only to verify that the conditions of Lemma 1.21 can be met in order to find an embedding of H into G. The idea is simple: we 'pad' H by adding $n - v(H)$ isolated vertices, and give a partition $\mathcal{X} = (X_i)_{i \in V(R)}$ of $V(H)$ in which X_i consists of $f^{-1}(i)$ together with isolated vertices such that $|V_i| = |X_i|$. We designate the isolated vertices in each X_i as the potential buffer vertices $\tilde{X}_i$. It then follows that the empty graph R' on $V(R)$ with no edges, together with these potential buffer vertices, give us a $(\gamma/2, R')$-buffer for $(H, \mathcal{X})$. By construction, the partition $\mathcal{X}$ is an R-partition of H. By definition, $\mathcal{V}$ is an (ε, d, p)-regular R-partition of G, and the super-regularity and inheritance properties required of R' are satisfied vacuously. Finally, we do not image restrict any vertices of H, so that the restriction pair properties are satisfied vacuously. Thus Lemma 1.21 gives us the desired embedding of H into G. □

We stress that the main difficulty in the proof of Theorem 1.8 is to obtain a spanning embedding; an almost-spanning embedding is much easier. It is similarly, and for similar reasons, much harder to prove Theorem 1.10 than Theorem 1.9. It is also worth noting that in the proof given in [28], there is a substantial amount of routine technical work to do in between obtaining size-compatible partitions of H and G and using the blow-up lemma to get an embedding of H into G, which is encapsulated in the so-called partial embedding lemma. This work is necessary because the blow-up lemma of [56] cannot be applied to the entire reduced graph. Our Lemma 1.21 can be applied to the entire reduced graph, and thus replaces both the partial embedding lemma and the blow-up lemma of [56].

6.5 Robustness of the Bandwidth Theorem

As with the proof of Theorem 1.9, the proof of Theorem 1.11 amounts to modifying the proof of Theorem 1.8. However, this time we need to rely on rather more of the machinery built up in [28]. Again, the reader not familiar with the argument there will wish to read this sketch in conjunction with Section 2 of [28].

Sketch proof of Theorem 1.11. We choose constants as in [28, Proof of Theorem 2], with the exception that we obtain ε' from Lemma 1.21 and ε from Theorem 2.7 and Corollary 2.8 for input ε' rather than from the blow-up lemma of [56] and the partial embedding lemma of [28]. We then follow the proof given there up to the point at which in [28] the first vertices of H are embedded using the partial embedding lemma. Let us recap what this amounts to. We are given graphs G and H satisfying the conditions of Theorem 1.11. We apply the Lemma for G of [28], which first returns a partition of $V(G)$ into parts $V'_{i,j}$, where $1 \leq i \leq k$ and $1 \leq j \leq \chi(H)$, with the following properties. Note that the Lemma for G does not explicitly return this partition, but it is convenient for the explanation to mention its existence; it also

does not explicitly give the upper bound on sizes of the parts, but this follows from the proof. First, $k \leq K_0 = K_0(\gamma, \chi(H))$ is bounded in terms of γ and $\chi(H)$. Second, the $(\varepsilon, d, 1)$-reduced graph R of this partition, whose vertex set is $[k] \times [\chi(H)]$ matching the indices of the partition, has minimum degree at least $\left(\frac{\chi(H)-1}{\chi(H)} + \frac{\gamma}{2}\right)v(R)$. Third, R contains a spanning backbone graph: that is, if $|i - i'| \leq 1$ and $j \neq j'$ then $(i, j)(i', j') \in E(R)$. Fourth, the parts $V'_{i,j}$ and $V'_{i,j'}$ differ in size by at most one, and each part has size between $(1 - \varepsilon)n/k\chi(H)$ and $2n/k\chi(H)$.

Now we apply the Lemma for H of [28], with the reduced graph R and the integer partition of n given by the $|V'_{i,j}|$. This gives us a homomorphism f from H to R, and a set of *special vertices* $Z \subseteq V(H)$, with the following properties, which depend on a quantity ξ satisfying $\beta \ll \xi \ll \varepsilon, K_0^{-1}$. First, $|Z| \leq k\chi(H)\xi n$. Second, for each $(i, j) \in V(R)$ we have $|f^{-1}(i, j)| = |V_{i,j}| \pm \xi n$. Third, if $uv \in H$ has neither endpoint in Z, then the first coordinates of $f(u)$ and $f(v)$ are equal, in other words u and v are mapped to vertices of the same clique in R. We set $X_{i,j} = f^{-1}(i, j)$ for each $(i, j) \in R$. By construction, the resulting partition $\mathcal{X}$ is an R-partition of H.

Next, we return to the Lemma for G, which guarantees, given the sizes of the parts $|X_{i,j}|$ satisfying the above properties, a partition $\mathcal{V}$ of $V(G)$ with parts $V_{i,j}$ for $(i, j) \in V(R)$ which has the following properties. First, $|V_{i,j}| = |X_{i,j}|$ for each $(i, j) \in V(R)$. Second, R is an $(\varepsilon, d, 1)$-reduced graph for G with respect to $\mathcal{V}$. Third, G is super-regular on the graph R' whose edges are $(i, j)(i, j')$ for $i \in [k]$ and $1 \leq j < j' \leq \chi(H)$, in other words on the $K_{\chi(H)}$-factor in the backbone graph.

This is the point at which, in [28], the embedding of H into G begins. It is worth remarking that, because Lemma 1.21 applies to the entire reduced graph, we could complete their proof by simply verifying the conditions of Lemma 1.21 for $p = 1$ (much as we do below), rather than needing the technical work of the partial embedding lemma.

Recall that we wish to show that H is a subgraph of G_p, where $p \geq C\left(\frac{\log n}{n}\right)^{1/\Delta}$ for some suitably large C. Observe that $E(G_p)$ is distributed as $E(G) \cap E(G(n, p))$. We now generate $\Gamma = G(n, p)$. Asymptotically almost surely, the good event of the blow-up lemma for random graphs, Lemma 1.21 occurs for Γ. We need to verify that H and $G_p = G \cap \Gamma$ a.a.s. satisfy the conditions of Lemma 1.21. We begin with H. Recall that $\mathcal{X}$ is an R-partition of H. By choice of ξ, it is 4-balanced. Also by choice of ξ, less than half of the vertices of any given $X_{i,j}$ are at distance two or less from Z. We let $\tilde{X}_{i,j}$ be the vertices of $X_{i,j}$ at distance three or more from Z. This gives us a $\left(\frac{1}{2}, R'\right)$-buffer for $(H, \mathcal{X})$.

We now need to show that a.a.s. $\mathcal{V}$ is an (ε', d, p)-regular R-partition of $G \cap \Gamma$, and that a.a.s. it is (ε', d, p)-super-regular and has one- and two-sided inheritance on R'. The first of these is an easy consequence of Theorem 2.1 (Chernoff's inequality) and the fact that $\mathcal{V}$ is an $(\varepsilon, d, 1)$-regular R-partition of G. Indeed, we can simply take the union bound over the at most 2^{2n} choices of pairs of subsets of $V(G)$ which we need to have density at least $(d - \varepsilon')p$. Since G has super-regularity on R', again using Theorem 2.1 and taking the union bound over the choices of vertices in $V(G)$ and $V(R')$, we see that a.a.s. $\Gamma \cap G$ has super-regularity on R'. Next we show that a.a.s. if $v \in V_i$ and $ij, jk \in R'$, so $\left(N_{\Gamma \cap G}(v; V_j), V_k\right)$ is (ε', d, p)-regular in G. Indeed, by Theorem 2.7, with $\beta = \frac{1}{4}$, the probability that this fails is at most 2^{-pn}, so that we can take a union bound over all choices of i, j, k and $v \in V_i$. Similarly, using Corollary 2.8, if also $ik \in R'$, a.a.s. $\left(N_{\Gamma \cap G}(v; V_j), N_{\Gamma \cap G}(v; V_k)\right)$ is (ε', d, p)-regular in G.

We do not image restrict any vertices of H, so the restriction pair condition of Lemma 1.21 is vacuously satisfied. Thus, applying Lemma 1.21, we see $H \subseteq G \cap \Gamma$ as desired. $\qquad\square$

Chapter 7

Concluding remarks

In this chapter we start by formulating a consequence of our random graphs blow-up lemma for dense graphs for later reference. We then discuss when our main results might be improved and when they are sharp (Section 7.2). We explain how to obtain (randomised) algorithmic versions of our blow-up lemmas (Section 7.3). We give a version of our random graphs blow-up lemma for directed graphs (Section 7.4) and sketch how one might allow for coloured graph settings, and mention hypergraphs (Section 7.5). Finally, we give some open problems (Section 7.6).

7.1 The dense case

One of the limitations of the dense blow-up lemma of [56] is that it can only be applied to small parts of the reduced graph of a partition given by the regularity lemma (because the regularity ε that the blow-up lemma requires depends on the number of clusters to which it is applied). Hence, when using this lemma to embed spanning graphs one usually applies the blow-up lemma several times to different small parts of the reduced graph after setting up suitable connections. In order to appropriately combine the connections with these blow-up lemma applications one uses image restrictions.

Our blow-up lemmas can be applied to the whole reduced graph. Since this can simplify blow-up lemma applications also in the dense case, we state a dense version here that is a direct consequence of Lemma 1.21. For easier reference we first repeat the relevant definitions needed for this lemma, which are simply the $p = 1$ case of the definitions introduced earlier. We remark that in particular our definition of restriction pairs simplifies to image restrictions (in the usual sense) in this case.

Let G, H, R, R' be graphs with $R' \subseteq R$. Let $\mathfrak{X} = \{X_i\}_{i \in [r]}$ be a partition of $V(H)$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$ be a partition of $V(G)$, let $\tilde{\mathfrak{X}} = \{\tilde{X}_i\}_{i \in [r]}$ be a family of subsets of $V(H)$, and let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a collection of subsets of $V(G)$.

- We say that the partition $\mathfrak{X}$ is κ-*balanced* if there exists $m \in \mathbb{N}$ such that we have $m \leq |X_i| \leq \kappa m$ for all $i, j \in [r]$. The partitions $\mathfrak{X}$ and $\mathcal{V}$ are *size-compatible* if $|X_i| = |V_i|$ for all $i \in [r]$.
- We say $(H, \mathfrak{X})$ is an *R-partition* if each part of X is nonempty, and whenever there are edges of H between X_i and X_j the pair ij is an edge of R.
- The family $\tilde{\mathfrak{X}} = \{\tilde{X}_i\}_{i \in [r]}$ is an (α, R')-*buffer* for $(H, \mathfrak{X})$ if for each $i \in [r]$ we have $\tilde{X}_i \subseteq X_i$ and

126

$|\tilde{X}_i| \geq \alpha|X_i|$, and for each $x \in \tilde{X}_i$ and $xy, yz \in E(H)$ with $x \in X_j$ and $z \in X_k$ we have $ij, jk \in E(R')$.

- $(G, \mathcal{V})$ is an (ε, d)-*regular R-partition* if for each $ij \in E(R)$ the pair (V_i, V_j) is (ε, d)-*regular*, that is, $d(V_i', V_j') \geq d - \varepsilon$ for all $V_i' \subseteq V_i$ with $|V_i'| \geq \varepsilon|V_i|$ and $V_j' \subseteq V_j$ with $|V_j'| \geq \varepsilon|V_j|$.

- $(G, \mathcal{V})$ is (ε, d)-*super-regular on R'* if for each $ij \in E(R')$ the pair (V_i, V_j) is (ε, d)-*super-regular*, that is, it is (ε, d)-regular and for every $u \in V_i$ we have $\deg_G(u; V_j) \geq (d - \varepsilon)|V_j|$ and for every $u \in V_j$ we have $\deg_G(u; V_i) \geq (d - \varepsilon)|V_i|$.

- We say that $\mathcal{I}$ is a family of (ρ, ζ)-*image restrictions* if the following properties hold for each $i \in [r]$ and $x \in X_i$.

  (a) The set $X_i^* \subseteq X_i$ of *image restricted* vertices in X_i, that is, vertices such that $I_x \neq V_i$, has size $|X_i^*| \leq \rho|X_i|$.

  (b) $|I_x| \geq \zeta|V_i|$.

Lemma 7.1 (Dense blow-up lemma for the whole reduced graph). *For all $\Delta, \Delta_{R'}, \kappa \geq 1$ and $\alpha, \zeta, d > 0$ there exists $\varepsilon, \rho > 0$ such that for all r_1 there is an n_0 such that for all $n \geq n_0$ the following holds. Let R be a graph on $r \leq r_1$ vertices and let $R' \subseteq R$ be a spanning subgraph with $\Delta(R') \leq \Delta_{R'}$. Let H and G be graphs with κ-balanced size-compatible vertex partitions $\mathcal{X} = \{X_i\}_{i \in [r]}$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$, respectively, which have parts of size at least $m \geq n/(\kappa r_1)$. Let $\tilde{\mathcal{X}} = \{\tilde{X}_i\}_{i \in [r]}$ be a family of subsets of $V(H)$ and let $\mathcal{I} = \{I_x\}_{x \in V(H)}$ be a collection of subsets of $V(G)$. Suppose that*

(B1) $\Delta(H) \leq \Delta$, $(H, \mathcal{X})$ is an R-partition, and $\tilde{\mathcal{X}}$ is an (α, R')-buffer for $(H, \mathcal{X})$,

(B2) $(G, \mathcal{V})$ is an (ε, d)-regular R-partition, which is (ε, d)-super-regular on R',

(B3) $\mathcal{I}$ is a family of (ρ, ζ)-image restrictions.

Then there is an embedding $\psi \colon V(H) \to V(G)$ such that $\psi(x) \in I_x$ for all $x \in V(H)$.

This lemma simplifies even further if image restrictions are not necessary, which is often the case in applications since this lemma can be applied to the reduced graph as a whole.

Note that the same lemma is proved as the $p = 1$ case of each of Lemmas 1.23 (D can be taken equal to 2Δ in this case, so that (*DBUL*4) becomes trivial) and 1.25. This is useful to note in that the algorithms used for the embedding are slightly different, and one might be easier to analyse than another for future extensions.

7.2 Optimality of our main results

In this section we discuss when our three blow-up lemmas can be improved. We expect that the statements of each lemma remain true for significantly smaller probabilities and weaker bijumbledness requirements in general, but we will observe that in some cases we cannot improve much, or to do so would require some additional conditions on the graphs into which we embed.

7.2.1 The blow-up lemma for random graphs

In the case $\Delta = 2$ and H contains a triangle, our result is optimal in terms of p up to the log factor, since it is well known that when $p = o(n^{-1/2})$ one can delete all triangles from $G(n, p)$ by removing only $o(pn^2)$ edges—so any 'blow-up-type' statement will be false.

When $\Delta = 3$, the statement of Lemma 1.21 is optimal up to the log factor. As we show in the next paragraph, in the event that H is a spanning K_4-factor, for small $c > 0$ and $p = cn^{-1/3}$, it is typically possible to find a subgraph G of $G(4n, p)$ with the following properties. There is a partition $V(G) = V_1 \cup \cdots \cup V_4$ with each set of size n. Letting $R = R' = K_4$, the partition is an $(\varepsilon, \frac{1}{2}, p)$-regular R-partition, which is super-regular, with one- and two-sided inheritance on R'. However, there is a vertex of G which is in no K_4, and thus there is no K_4-factor covering G.

To construct G, we let $V_1, \ldots, V_4$ be an equipartition of $[4n]$, fix a vertex $v \in V_1$, reveal $G(4n, p)$, remove all edges within each part, and then remove from $N_G(v)$ all triangles by deleting a minimum number of edges. It is not hard to check using the Chernoff bound, Theorem 2.1, that a.a.s. before this final step, we have an $(\varepsilon, \frac{3}{4}, p)$-regular R-partition with super-regularity and one- and two-sided inheritance on R'. It is similarly easy to check that a.a.s. before the final step the degree of each vertex in V_i to any other V_j is close to pn, and any pair of vertices has about $p^2 n$ common neighbours in other parts. Thus the expected number of triangles in $N_G(v)$ is $p^3 (pn)^3 = c^6 n$, and the actual number is somewhat concentrated, so that a.a.s. in the final step we delete at most $2c^6 n$ edges. This is far too small to destroy one-sided inheritance, and, for $c > 0$ sufficiently small, too small to destroy two-sided inheritance. It remains to show that we have not removed too many edges from any one vertex. But the expected number of edges removed from $u \in N_G(v)$ is at most $p(p^2 n)^2 = c^3 p^2 n$; the actual number is by Theorem 2.1 exponentially concentrated, and hence a.a.s. we remove $o(pn)$ edges from each u, so do not destroy super-regularity.

Perhaps inserting an extra condition into the statement, such as insisting that all vertex neighbourhoods contain $\Omega(p^6 n^3)$ triangles, would allow one to prove a blow-up lemma which allows for embedding a K_4-factor down to the natural limit $p = n^{-2/5}$. However, it would not always be possible to obtain such a condition. In [6] it is shown that Breaker wins the K_4-factor game on K_n with a bias $13n^{1/3}$ (Lemma 1.21 is used to show that Maker wins when the bias is $c\left(\frac{n}{\log n}\right)^{1/3}$ for some small $c > 0$). It follows that any 'extra condition' is one which Maker cannot guarantee to obtain with bias $13n^{1/3}$.

Finally, for $\Delta \geq 4$ one might hope that the statement of our blow-up lemma remains true down to $p = Cn^{-2/(\Delta+3)}$, this being the point at which one can generalise the above construction (and therefore the point at which the statement provably does not guarantee a $K_{\Delta+1}$-factor). Perhaps, optimistically, one might hope that there are some natural extra conditions which even allow $p \geq Cn^{-2/(\Delta+2)}$, this being the point at which one can remove all copies of $K_{\Delta+1}$ by deleting a tiny fraction of the edges of $G(n, p)$.

However, we believe that improving upon our result is likely to be very challenging. Even in the (much simpler and well-studied) setting of trying to prove $\mathcal{H}(n, \Delta)$-universality of $G(n, p)$, until very recently no improvement had been made over what one can obtain from Lemma 1.21. Very recently Ferber and Nenadov [39] were able to obtain a small improvement. In the almost-spanning setting, Conlon, Ferber, Nenadov and Škorić [30] could improve more on what follows from Lemma 1.21, showing that the random graph $G((1+\gamma)n, p)$ is a.a.s. $\mathcal{H}(n, \Delta)$-universal when $p = \omega(n^{-1/(\Delta-1)} \log^5 n)$ for $\Delta \geq 3$. But the improvement in the exponent one would desire is of order Δ^{-1} not Δ^{-2}.

7.2.2 The blow-up lemma for bijumbled graphs

We do not believe the $(p, cp^{\max(4,(3\Delta+1)/2)})$-bijumbledness requirement in Lemma 1.25 is optimal. Most of the proof would work with $(p, cp^{\Delta+2}n)$-bijumbledness, but we were not able to find a way to avoid the LNS property which requires the stronger condition. Nevertheless we conjecture that $(p, cp^{\Delta+1}n)$-

bijumbledness suffices (we expect that the extra factor of p can be gained by using reserved cliques as in the proof of Lemma 1.21). It is still not clear that this would be optimal. For $\Delta = 2$ we need $(p, cp^2 n)$-bijumbledness, since Alon [9] constructed a $(p, Cp^2 n)$-bijumbled graph without triangles. It is a believable conjecture (see for example Conlon, Fox and Zhao [31]) that there are $(p, Cp^\Delta n)$-bijumbled graphs without $K_{\Delta+1}$ for every Δ, in which case the same requirement would be necessary for a blow-up lemma. It is possible, however, (conjectured in [31], but the contrary is conjectured by Kohayakawa, Rödl, Schacht and Skokan [53]) that copies of $K_{\Delta+1}$ cannot be guaranteed in regular subgraphs of (p, β)-bijumbled graphs at this point, but instead that $\beta \le cp^{\Delta+1} n$ is required (it is proved in [31] that at this point one can guarantee copies of $K_{\Delta+1}$).

Note that Lemma 1.25 does not permit linearly many image restrictions. The reason for this is that our method for queue embedding requires that underlying restriction sets tend not to cluster, which is encapsulated in (4.6) of Claim 4.5. But we can only establish the inequality (4.6) by excluding image restricted vertices (and in fact the inequality would not necessarily hold if linearly many image restricted vertices were included). We can only image restrict about a p^Δ-fraction of vertices in each cluster, as these are too few to cause problems with the queue embedding. In many applications this is not a problem (see for example [3]), but it could well cause a problem for some applications. We believe that it is possible to modify the proof strategy substantially in order to have some control over linearly many vertices of each X_i. The modification we would make is the following. We would permit the user of the blow-up lemma to specify $\rho|X_i|$ 'pre-embedded' vertices in each part X_i which are to be embedded first (before even the neighbours of buffer vertices). The user is then permitted to embed these vertices sequentially, subject to four conditions. First, the result must be a good partial embedding ψ. Second, each vertex must be embedded to a uniform random vertex from a set of size at least $\frac{1}{10}\mu\zeta(dp)^{\Delta-1} n/(\kappa r_1)$. Third, for each $jk \in R'$ and $v \in V_j$, at most $\frac{1}{100}\deg_G(v; V_k)$ neighbours of v in V_k may be in $\mathrm{Im}(\psi)$. Fourth, the total number of vertices in X_i with pre-embedded neighbours is at most $\rho|X_i|$. We would then follow the proof strategy of Lemma 1.25 to embed the remainder of H into G, treating the neighbours of pre-embedded vertices as image restricted. We remark that this does not automatically resolve the problem with Claim 4.5, since the partition of $V(G)$ considered at this point is *not* the same as the partition the user of the blow-up lemma supplies. But it is easy to modify Lemma 2.22 (the good partitions lemma, which generates the partition considered by Claim 4.5) to show that all the sets under consideration are with high probability evenly distributed by the random equipartition, and then the proof of Claim 4.5 does go through. Checking the full details of this approach, and for that matter using the resulting blow-up lemma, seems likely to be non-trivial, but it could potentially allow for stronger theorems.

7.2.3 The blow-up lemma to embed degenerate graphs

The definition of (D, p, m)-boundedness is rather complicated, but it allows us to prove as flexible a statement as we could, taking into account the structure of H in order to work with lower probabilities. Observe that a d-degenerate n-vertex graph with bounded maximum degree, which Lemma 1.23 can handle with $p \approx n^{-1/(2d+1)}$, can contain almost dn edges. This is comparable to a $2d$-regular graph for which Lemma 1.21 would require $p \approx n^{-1/2d}$. So even without saying anything about the structure of H beyond its degeneracy, the result is almost as powerful as Lemma 1.21.

In the event that we only need an almost-spanning embedding, we can take the potential buffer vertices in each part to be isolated, and hence $D = 2d$. We can then embed d-degenerate graphs with $p \approx n^{-1/(2d)}$,

matching the performance of Lemma 1.21. Finally, if H is an F-factor we can take $D = d + 3$ or better (depending on the structure of F), in which case the performance of Lemma 1.23 substantially improves, working with $p \approx n^{-1/(d+3)}$.

In fact, we can improve Lemma 1.23 for F-factors. As stated in the proof of Lemma 5.2 (the degenerate RGA lemma) we require (*ORD* 3) and the condition $\pi^\tau(x) \leq D_x - 1$ for vertices x adjacent to potential buffer vertices within (*ORD* 1) only in order to prove (*DRGA* 3). When embedding an F-factor (provided R' is suitable, for example $R' = K_{|F|}$), we do not really need (*DRGA* 3), as (*DRGA* 2) shows that only a few vertices in each cluster of G can fail (*DRGA* 3), and we can use an argument similar to Lemma 3.3 to deal with them. This allows us to reduce the required D by one, compared to the requirement of Lemma 1.23.

In particular, for each $s \leq t$ one can prove a blow-up lemma which embeds a $K_{s,t}$-factor when $p \geq C\left(\frac{\log n}{n}\right)^{1/s}$. This is almost optimal in terms of p, since the 2-density of $K_{s,t}$ is $\frac{st-1}{s+t-2}$, which approaches s as t becomes large; when p is below $n^{-(s+t-2)/(st-1)}$, one can delete a very small fraction of the edges of $G(n, p)$ to destroy all copies of $K_{s,t}$.

7.3 Algorithmic embedding

The proofs of our blow-up lemmas can be changed slightly to give polynomial-time randomised algorithms which with high probability construct the embeddings we prove exist. It is quite tedious to check the details, but we provide a sketch of this for the interested reader.

The main change which needs to be made concerns the certification of sparse-regular pairs. The RI property (see Section 2.2), which we require in all of our blow-up lemmas, guarantees that typical vertex neighbourhoods inherit regularity, but for an algorithm we need to be able to identify the set of vertices whose neighbourhoods inherit regularity in polynomial time.

Alon, Duke, Lefmann, Rödl and Yuster [15] showed that, in dense graphs, determining if a given bipartite graph is ε-fully-regular is co-NP-complete, but that there is a polynomial-time algorithm which either certifies ε-full-regularity, or returns a witness to the failure of ε'-full-regularity, for some ε' which may be much smaller than ε but does not depend on the number of vertices in the regular pair. For sparse graphs, a corresponding polynomial-time certification algorithm was given by Alon, Coja-Oghlan, Hàn, Kang, Rödl and Schacht [14] which either certifies (ε, d, p)-full-regularity or returns a witness to the failure of (ε', d, p)-full-regularity.

Unfortunately, we do not know of any such algorithm in the literature for lower-regular pairs, so we now sketch a certification algorithm, using the results of [14], for lower-regularity which works in subgraphs of random graphs. Given as input ε, d, p and a bipartite graph, which must be bounded (the definition is in [14], but the reader does not need to know it), we apply the algorithmic sparse regularity lemma of [14] with regularity parameter $\frac{1}{100}\varepsilon^3$. This returns a partition of each side of the bipartite graph into approximately equal numbers of parts, which approximately equipartition each side, even if the bipartite graph itself is very unbalanced (we may want, for example, to know whether a bipartite graph with parts of size pn and p^2n respectively is lower-regular). We choose $0 < \varepsilon' < \frac{1}{100}\varepsilon^3$ such that any given part of the partition contains at least an ε'-fraction of its side. Now if any pair of parts in this partition has density less than $(d - \varepsilon')p$, it is a witness to a failure of (ε', d, p)-lower-regularity. If not, we claim the bipartite graph is (ε, d, p)-lower-regular; the choice of regularity parameter in the use of

sparse regularity ensures that there are too few irregular pairs to seriously affect densities between large sets. We note that the requirement of boundedness is needed for the algorithmic sparse regularity lemma, and that this boundedness holds a.a.s. in subgraphs of typical random graphs (and is implied by the NS property which we require in any case). In contrast to the certification algorithm for sparse full-regularity, the dependency of ε' on ε here is very poor: there is a tower-type relationship, which appears iterated in the constant dependencies of the algorithmic versions of Lemmas 1.21 and 1.23.

For either random or bijumbled ambient graphs Γ, given a certification algorithm and inheritance lemmas, one can prove a variant of the RI property in which not only do typical vertex neighbourhoods inherit (either version of) regularity, but they do so certifiably. We follow the proofs more or less as in Section 2.2, except that at each step, where we need certifiable (ε, d, p)-regularity for some ε, we obtain ε' from the certification algorithm and then let ε'' be returned by our inheritance lemmas for input ε'. Now a typical vertex neighbourhood is (ε', d, p)-regular, so that the certification algorithm will certify it to be (ε, d, p)-regular as desired.

We now sketch how one can use this to obtain algorithmic versions of our blow-up lemmas.

It is necessary to check that the algorithm implicit in the proof of the lemma obtaining good partitions of H and G (Lemma 2.22) is a randomised polynomial time algorithm. This follows since the Hajnal-Szemerédi theorem (Theorem 2.3) has an algorithmic version [48], and since the proof of Lemma 2.4 can then easily be made constructive. Furthermore, a failure of the randomised partitioning to produce a good G-partition can be detected in polynomial time.

It is further necessary to check that each of our RGA algorithms (Algorithms 1, 3 and 4) can be carried out in polynomial time. This amounts to checking that the various sets that appear in these algorithms can be constructed in polynomial time. For most of these sets, this is obviously possible. However, in order to construct the bad sets $B_t(x)$ we need the certifiable regularity inheritance discussed above.

In proving Lemma 1.21 we give (implicitly) an algorithm for embedding queue vertices. Again, to run this algorithm we need to construct the bad sets $B(x)$ and this requires certifiable regularity inheritance. We also use the fact that the bipartite matching problem can be solved in polynomial time (see Kuhn [66]). For fixing buffer defects in polynomial time, Lemma 3.7 gives sets P_i and D_i, and the proof constructs these sets in polynomial time, while our algorithm for fixing buffer defects, Algorithm 2, which uses the P_i and D_i, is trivially polynomial time.

Finally, each of our blow-up lemmas is completed by embedding the buffer vertices. This amounts to a bipartite matching problem, and can be solved in polynomial time.

It seems reasonable to believe that it is possible to derandomise our RGA algorithms and good partitions algorithm, which would yield polynomial time algorithms for constructing each of the claimed embeddings. Certainly Komlós, Sárközy and Szemerédi [57] were able to derandomise their original (RGA-based) proof of the dense blow-up lemma. However, we did not attempt to check whether their methods suffice in our case.

7.4 Directed graphs

Although our blow-up lemmas as written apply to undirected graphs, we can also apply them to subdigraphs of random directed graphs (or bijumbled directed graphs). The random directed graph $\overrightarrow{G}(n,p)$ is obtained by choosing, for each ordered pair of vertices in a vertex set of size n, to put an arc independently

with probability p. We give for illustration the directed statement corresponding to Lemma 1.21. In order to state this, we define the *undirection* of a digraph $\overrightarrow{G}$ to be the graph G with $uv \in E(G)$ if and only if either $\overrightarrow{uv}$ or $\overrightarrow{vu}$ is an arc of $\overrightarrow{G}$. The terms which we defined for undirected graphs (such as R-partition, buffer, and so on) are taken as applying to the undirections of the digraphs considered. The exception is that when we talk about an $\overrightarrow{R}$-partition of $\overrightarrow{G}$ or $\overrightarrow{H}$ we require, in addition to the conditions for the undirections of digraphs, that all arcs of $\overrightarrow{G}$ or $\overrightarrow{H}$ go in the direction specified by the arcs of $\overrightarrow{R}$. In the following lemma we work with $\overrightarrow{G}(n,q)$ but use $p = 2q - q^2$ as the parameter for regularity. The reason for this is that the undirection of $\overrightarrow{G}(n,q)$ is $G(n,p)$.

Lemma 7.2 (Blow-up lemma for random directed graphs). *For all Δ, $\Delta_{R'}$, Δ_J, $\alpha, \zeta, d > 0$, $\kappa > 1$ there exist $\varepsilon, \rho > 0$ such that for all r_1 there is a C such that for $q > C(\log n/n)^{1/\Delta}$ the following holds. Let $p = 2q - q^2$. The random directed graph $\overrightarrow{\Gamma} = \overrightarrow{G}(n,q)$ asymptotically almost surely satisfies the following.*

Let $\overrightarrow{R}$ be a digraph on $r \le r_1$ vertices without cycles of length 2 and let R' be a subgraph of the undirection of $\overrightarrow{R}$ with $\Delta(R') \le \Delta_{R'}$. Let $\overrightarrow{H}$ and $\overrightarrow{G} \subseteq \overrightarrow{\Gamma}$ be digraphs given with κ-balanced, size-compatible vertex partitions $\mathfrak{X} = \{X_i\}_{i \in [r]}$ and $\mathcal{V} = \{V_i\}_{i \in [r]}$ with parts of size at least $m \ge n/(\kappa r_1)$. Let $\mathfrak{I} = \{I_x\}_{x \in V(\overrightarrow{H})}$ be a family of image restrictions, and $\mathfrak{J} = \{J_x\}_{x \in V(\overrightarrow{H})}$ be a family of restricting vertices. Suppose that

(BUL 1) *The undirection of $\overrightarrow{H}$ has maximum degree at most Δ, $(\overrightarrow{H}, \mathfrak{X})$ is an $\overrightarrow{R}$-partition, and $\tilde{\mathfrak{X}} = \{\tilde{X}_i\}_{i \in [r]}$ is an (α, R')-buffer for $\overrightarrow{H}$,*

(BUL 2) *$(\overrightarrow{G}, \mathcal{V})$ is an (ε, d, p)-regular $\overrightarrow{R}$-partition, which is (ε, d, p)-super-regular on R', has one-sided inheritance on R', and two-sided inheritance on R' for $\tilde{\mathfrak{X}}$,*

(BUL 3) *$\mathfrak{I}$ and $\mathfrak{J}$ form a $(\rho, \zeta, \Delta, \Delta_J)$-restriction pair.*

Then there is an embedding $\psi \colon V(\overrightarrow{H}) \to V(\overrightarrow{G})$ such that $\psi(x) \in I_x$ for each $x \in \overrightarrow{H}$.

This lemma is a corollary of Lemma 1.21. We simply work with the undirections of the given digraphs, which satisfy the conditions of Lemma 1.21, and observe that an embedding ψ of the undirection of $\overrightarrow{H}$ into the undirection of $\overrightarrow{G}$ such that $\psi(x) \in V_i$ for each $x \in X_i$ is by definition of an $\overrightarrow{R}$-partition automatically an embedding of $\overrightarrow{H}$ into $\overrightarrow{G}$.

With rather more work, we believe we could allow $\overrightarrow{R}$ to contain 2-cycles. More generally, we could allow G to be coloured from a palette of at most κ colours, define R to be the multicoloured graph (with edges permitted to have several colours) corresponding to relatively dense sparse-regular pairs in the given colour, and supply a coloured graph H with a coloured graph homomorphism to R, which we require to be embedded with edges going to edges of G with the correct colour. We believe that such coloured versions of all three of our blow-up lemmas can be proved, following the strategies given in this paper. However, to do so requires appropriate modifications to several definitions and recalculation of various parameters. We see no reason why this should cause difficulty, but we did not check the details.

7.5 Hypergraphs

It seems likely that the techniques developed in this paper will be very helpful for proving a blow-up lemma for uniform hypergraphs which works relative to sparse random or pseudorandom (appropriately

defined) hypergraphs. In the dense case, Keevash [47] proved a hypergraph blow-up lemma. However, it has a serious limitation, in that it only allows for image restriction of single vertices. In many applications one needs image restrictions of vertex tuples of size up to $k-1$ in k-uniform hypergraphs.

7.6 Open problems

Beyond the question of improving on our main results (as discussed in Section 7.2), we would like to pose the following problems.

Problem 7.3. Is it true that for each $r, \Delta \geq 2$ and n there exists a $K_{\Delta+1,\Delta+1}$-free graph G which is r-partition universal for $\mathcal{H}(n,\Delta)$, with $v(G) = O(n)$?

We have seen that the answer is 'yes' if $\Delta + 1$ is replaced by 2Δ (Theorem 1.5), and trivially the answer is 'no' if $\Delta + 1$ is replaced by Δ. It is not clear that G should be constructed randomly (which is how Theorem 1.5 is proved) in order to obtain an affirmative answer to this problem. But it is also not clear how to construct sparse graphs with strong Ramsey properties any other way.

Before trying to improve the random graph blow-up lemmas in this paper, Lemmas 1.21 and 1.23, we should at least know how to embed large subgraphs in the random graph itself in a robust way.

Problem 7.4. For what p is $G(n,p)$ typically $\mathcal{H}(n,\Delta)$-universal, or typically $\mathcal{H}(n,d,\Delta)$-universal? Does the answer change substantially if $G(n,p)$ is replaced with $G(Cn,p)$ for C large, or for C slightly larger than one?

So far, all the methods used to attack problems of this type construct a single embedding of some H into $G(n,p)$ step by step, and rely on there being many ways to continue the embedding after each step. In this paper, each step is a single vertex embedding. In order for there to be many ways to embed a vertex with Δ embedded neighbours, we certainly need $p^\Delta n \gg 1$, and Lemma 1.21 gets within a log-factor of this bound. Recent papers, in particular [30] and [39], improve slightly on this bound by embedding more than one vertex at a time. But it seems reasonable to conjecture that if $p \geq Cn^{-2/(\Delta+1)} \log^{1/\binom{\Delta+1}{2}}$ then $G(n,p)$ should typically be $\mathcal{H}(n,\Delta)$-universal. For large Δ, this is much sparser than anything we can currently handle, and it seems likely that being more clever with the step-by-step embedding methods currently used will not prove this conjecture. Assuming the conjecture is correct, it can make at most a log-factor difference if $G(n,p)$ is replaced by $G(Cn,p)$ for C slightly larger than 1, or for C large. However, it should be significantly easier to prove universality in the latter two cases.

For $\mathcal{H}(n,d,\Delta)$-universality of $G(Cn,p)$ with C slightly larger than 1, an almost optimal bound on p is obtained in [69, Theorem 3.6]. This problem is easier: it is necessary for $p^d n$ to be large in order for a vertex-by-vertex embedding method to work, but it is easy to see that this is also a lower bound for $\mathcal{H}(n,d,\Delta)$-universality. It would still be nice to remove the extra log-factors from [69, Theorem 3.6]. What seems much harder is to prove a similar spanning universality result.

Problem 7.5. For what β are (p,β)-bijumbled n-vertex graphs G with minimum degree $\frac{1}{2}pn$ always $\mathcal{H}(n,\Delta)$-universal? Does the answer change if we allow G to have Cn vertices for C large, or for C slightly larger than one?

This problem is a first step towards improving Lemma 1.25. It would also be interesting to replace bijumbledness with one of the other standard notions of pseudorandomness. However, in this setting we do not really know how large β can be: we do not have good constructions of graphs which are strongly bijumbled but are not $\mathcal{H}(n,\Delta)$-universal, except for Alon's construction [9] for $\Delta = 2$.

Problem 7.6. Do there exist graphs G which are r-partition universal for $\mathcal{H}(n,\Delta)$ with only $O(n^{2-2/\Delta})$ edges?

Alon and Capalbo [11] showed that for $r = 1$ (i.e. just universality) the answer is yes and that this is best possible. This is the best known lower bound for larger r, and it would be interesting to know if it is correct. Here a simple random construction (which is used in [53] to obtain the best known upper bound on the number of edges required) cannot work: a random graph with this many edges will not contain a $K_{\Delta+1}$-factor at all, let alone have the Ramsey property for it.

Problem 7.7. For what bias b can Maker win the $\mathcal{H}(n,\Delta)$-universality game on K_n, or on K_{Cn} for C large, or for C slightly larger than one?

Although there exist universal graphs with $n^{2-2/\Delta}$ edges (Alon and Capalbo [11]), Maker certainly cannot make them with a bias $\Omega(n^{2/\Delta})$, since Maker requires $b = O(n^{2/(\Delta+2)})$ in order to make just one copy of $K_{\Delta+1}$.

Bibliography

[1] Peter Allen and Julia Böttcher, *Partition universality for graphs of bounded degeneracy and degree*. 8

[2] Peter Allen, Julia Böttcher, Julia Ehrenmüller, Jakob Schnitzer, and Anusch Taraz, *A spanning bandwidth theorem in random graphs*, Combin. Probab. Comput. **31** (2022), no. 4, 598–628. MR 4439774 11

[3] Peter Allen, Julia Böttcher, Julia Ehrenmüller, and Anusch Taraz, *The bandwidth theorem in sparse graphs*, Adv. Comb. (2020), Paper No. 6, 60. MR 4125347 10, 11, 15, 19, 20, 22, 129

[4] Peter Allen, Julia Böttcher, Simon Griffiths, Yoshiharu Kohayakawa, and Robert Morris, *Chromatic thresholds in dense random graphs*, Random Structures Algorithms **51** (2017), no. 2, 185–214. MR 3683361 22

[5] Peter Allen, Julia Böttcher, Hiệp Hàn, Yoshiharu Kohayakawa, and Yury Person, *Powers of Hamilton cycles in pseudorandom graphs*, Combinatorica **37** (2017), no. 4, 573–616. MR 3694704 7, 21

[6] Peter Allen, Julia Böttcher, Yoshiharu Kohayakawa, Humberto Naves, and Yury Person, *Making spanning graphs*. 9, 128

[7] Peter Allen, Julia Böttcher, Yoshiharu Kohayakawa, and Barnaby Roberts, *Triangle-free subgraphs of random graphs*, Combin. Probab. Comput. **27** (2018), no. 2, 141–161. MR 3778197 22

[8] Peter Allen, Julia Böttcher, Jozef Skokan, and Maya Stein, *Regularity inheritance in pseudorandom graphs*, Random Structures Algorithms **56** (2020), no. 2, 306–338. MR 4060348 22, 30, 31

[9] Noga Alon, *Explicit Ramsey graphs and orthonormal labelings*, Electron. J. Combin. **1** (1994), Research Paper 12, 8pp. MR 1302331 129, 134

[10] Noga Alon and Jean Bourgain, *Additive patterns in multiplicative subgroups*, Geom. Funct. Anal. **24** (2014), no. 3, 721–739. MR 3213827 7

[11] Noga Alon and Michael Capalbo, *Sparse universal graphs for bounded-degree graphs*, Random Structures Algorithms **31** (2007), no. 2, 123–133. MR 2343715 7, 134

[12] _____ , *Optimal universal graphs with deterministic embedding*, Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, 2008, pp. 373–378. MR 2485323 7

[13] Noga Alon, Michael Capalbo, Yoshiharu Kohayakawa, Vojtěch Rödl, Andrzej Ruciński, and Endre Szemerédi, *Universality and tolerance (extended abstract)*, 41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000), IEEE Comput. Soc. Press, Los Alamitos, CA, 2000, pp. 14–21. MR 1931800 5

[14] Noga Alon, Amin Coja-Oghlan, Hiêp Hàn, Mihyun Kang, Vojtěch Rödl, and Mathias Schacht, *Quasi-randomness and algorithmic regularity for graphs with general degree distributions*, SIAM J. Comput. **39** (2010), no. 6, 2336–2362. MR 2644348 130

[15] Noga Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster, *The algorithmic aspects of the regularity lemma*, J. Algorithms **16** (1994), no. 1, 80–109. MR 1251840 130

[16] Noga Alon and Zoltán Füredi, *Spanning subgraphs of random graphs*, Graphs Combin. **8** (1992), no. 1, 91–94. MR 1157513 23

[17] Noga Alon, Michael Krivelevich, Joel Spencer, and Tibor Szabó, *Discrepancy games*, Electron. J. Combin. **12** (2005), Research Paper 51, 9pp. MR 2176527 9

[18] Noga Alon and Joel H. Spencer, *The probabilistic method*, third ed., Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2008, With an appendix on the life and work of Paul Erdős. MR 2437651 6

[19] József Balogh, Choongbum Lee, and Wojciech Samotij, *Corrádi and Hajnal's theorem for sparse random graphs*, Combin. Probab. Comput. **21** (2012), no. 1-2, 23–55. MR 2900047 5

[20] József Balogh, Robert Morris, and Wojciech Samotij, *Independent sets in hypergraphs*, J. Amer. Math. Soc. **28** (2015), no. 3, 669–709. MR 3327533 5

[21] Victor Bapst and Amin Coja-Oghlan, *Harnessing the Bethe free energy*, Random Structures Algorithms **49** (2016), no. 4, 694–741. MR 3570985 4

[22] József Beck, *Combinatorial games*, Encyclopedia of Mathematics and its Applications, vol. 114, Cambridge University Press, Cambridge, 2008, Tic-tac-toe theory. MR 2402857 9

[23] Małgorzata Bednarska and Tomasz Łuczak, *Biased positional games for which random strategies are nearly optimal*, Combinatorica **20** (2000), no. 4, 477–488. MR 1804821 9

[24] Béla Bollobás, *Random graphs*, second ed., Cambridge Studies in Advanced Mathematics, vol. 73, Cambridge University Press, Cambridge, 2001. MR 1864966 5

[25] Julia Böttcher, Yoshiharu Kohayakawa, and Anusch Taraz, *Almost spanning subgraphs of random graphs after adversarial edge removal*, Combin. Probab. Comput. **22** (2013), no. 5, 639–683. MR 3094477 5, 10, 122, 123

[26] Julia Böttcher, Yoshiharu Kohayakawa, Anusch Taraz, and Andreas Würfl, *An extension of the blow-up lemma to arrangeable graphs*, SIAM J. Discrete Math. **29** (2015), no. 2, 962–1001. MR 3353133 23

[27] Julia Böttcher, Klaas P. Pruessmann, Anusch Taraz, and Andreas Würfl, *Bandwidth, expansion, treewidth, separators and universality for bounded-degree graphs*, European J. Combin. **31** (2010), no. 5, 1217–1227. MR 2644412 10

[28] Julia Böttcher, Mathias Schacht, and Anusch Taraz, *Proof of the bandwidth conjecture of Bollobás and Komlós*, Math. Ann. **343** (2009), no. 1, 175–205. MR 2448444 10, 123, 124, 125

[29] Amin Coja-Oghlan, Will Perkins, and Kathrin Skubch, *Limits of discrete distributions and Gibbs measures on random graphs*, European J. Combin. **66** (2017), 37–59. MR 3692136 4

[30] David Conlon, Asaf Ferber, Rajko Nenadov, and Nemanja Škorić, *Almost-spanning universality in random graphs*, Random Structures Algorithms **50** (2017), no. 3, 380–393. MR 3632416 7, 128, 133

[31] David Conlon, Jacob Fox, and Yufei Zhao, *Extremal results in sparse pseudorandom graphs*, Adv. Math. **256** (2014), 206–290. MR 3177293 5, 6, 22, 129

[32] David Conlon, W. T. Gowers, W. Samotij, and M. Schacht, *On the KłR conjecture in random graphs*, Israel J. Math. **203** (2014), no. 1, 535–580. MR 3273450 5

[33] David Conlon, Rajko Nenadov, and Miloš Trujić, *The size-Ramsey number of cubic graphs*, Bull. Lond. Math. Soc. **54** (2022), no. 6, 2135–2150. MR 4528614 8

[34] Domingos Dellamonica, Jr., Yoshiharu Kohayakawa, Martin Marciniszyn, and Angelika Steger, *On the resilience of long cycles in random graphs*, Electron. J. Combin. **15** (2008), no. 1, Research Paper 32, 26pp. MR 2383452 21

[35] Domingos Dellamonica, Jr., Yoshiharu Kohayakawa, Vojtěch Rödl, and Andrzej Ruciński, *An improved upper bound on the density of universal random graphs*, Random Structures Algorithms **46** (2015), no. 2, 274–299. MR 3302898 6, 14

[36] Nemanja Draganić and Kalina Petrova, *Size-Ramsey numbers of graphs with maximum degree three*. 8

[37] Paul Erdős and George Szekeres, *A combinatorial problem in geometry*, Compositio Math. **2** (1935), 463–470. MR 1556929 120

[38] Asaf Ferber, Michael Krivelevich, and Humberto Naves, *Generating random graphs in biased maker-breaker games*, Random Structures Algorithms **47** (2015), no. 4, 615–634. MR 3418908 9, 15, 121

[39] Asaf Ferber and Rajko Nenadov, *Spanning universality in random graphs*, Random Structures Algorithms **53** (2018), no. 4, 604–637. MR 3875277 6, 128, 133

[40] Asaf Ferber, Rajko Nenadov, and Ueli Peter, *Universality of random graphs and rainbow embedding*, Random Structures Algorithms **48** (2016), no. 3, 546–564. MR 3481273 6

[41] Jon Folkman, *Graphs with monochromatic complete subgraphs in every edge coloring*, SIAM J. Appl. Math. **18** (1970), 19–24. MR 268080 8

[42] Stefanie Gerke, Yoshiharu Kohayakawa, Vojtěch Rödl, and Angelika Steger, *Small subsets inherit sparse ε-regularity*, J. Combin. Theory Ser. B **97** (2007), no. 1, 34–56. MR 2278123 13, 22, 31, 32

[43] Ben Green and Terence Tao, *The primes contain arbitrarily long arithmetic progressions*, Ann. of Math. (2) **167** (2008), no. 2, 481–547. MR 2415379 4

[44] András Hajnal and Endre Szemerédi, *Proof of a conjecture of P. Erdős*, Combinatorial theory and its applications, I-III (Proc. Colloq., Balatonfüred, 1969), Colloq. Math. Soc. János Bolyai, vol. 4, North-Holland, Amsterdam-London, 1970, pp. 601–623. MR 297607 29

[45] Hao Huang, Choongbum Lee, and Benny Sudakov, *Bandwidth theorem for random graphs*, J. Combin. Theory Ser. B **102** (2012), no. 1, 14–37. MR 2871764 10, 38

[46] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000. MR 1782847 5, 26

[47] Peter Keevash, *A hypergraph blow-up lemma*, Random Structures Algorithms **39** (2011), no. 3, 275–376. MR 2816936 133

[48] Henry A. Kierstead, Alexandr V. Kostochka, Marcelo Mydlarz, and Endre Szemerédi, *A fast algorithm for equitable coloring*, Combinatorica **30** (2010), no. 2, 217–224. MR 2676836 131

[49] Jeong Han Kim and Sang June Lee, *Universality of random graphs for graphs of maximum degree two*, SIAM J. Discrete Math. **28** (2014), no. 3, 1467–1478. MR 3259784 6

[50] Yoshiharu Kohayakawa, *Szemerédi's regularity lemma for sparse graphs*, Foundations of computational mathematics (Rio de Janeiro, 1997), Springer, Berlin, 1997, pp. 216–230. MR 1661982 4, 120

[51] Yoshiharu Kohayakawa, Tomasz Łuczak, and Vojtěch Rödl, *On K^4-free subgraphs of random graphs*, Combinatorica **17** (1997), no. 2, 173–213. MR 1479298 4

[52] Yoshiharu Kohayakawa and Vojtěch Rödl, *Regular pairs in sparse random graphs. I*, Random Structures Algorithms **22** (2003), no. 4, 359–434. MR 1980964 13

[53] Yoshiharu Kohayakawa, Vojtěch Rödl, Mathias Schacht, and Jozef Skokan, *On the triangle removal lemma for subgraphs of sparse pseudorandom graphs*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 359–404. MR 2815608 129, 134

[54] Yoshiharu Kohayakawa, Vojtěch Rödl, Mathias Schacht, and Endre Szemerédi, *Sparse partition universal graphs for graphs of bounded degree*, Adv. Math. **226** (2011), no. 6, 5041–5065. MR 2775894 5, 8, 12, 13, 22, 23, 24, 31, 39, 40, 121

[55] János Komlós, *The blow-up lemma*, vol. 8, 1999, Recent trends in combinatorics (Mátraháza, 1995), pp. 161–176. MR 1684627 4

[56] János Komlós, Gábor N. Sárközy, and Endre Szemerédi, *Blow-up lemma*, Combinatorica **17** (1997), no. 1, 109–123. MR 1466579 4, 7, 12, 15, 16, 17, 23, 24, 60, 62, 124, 126

[57] ———, *An algorithmic version of the blow-up lemma*, Random Structures Algorithms **12** (1998), no. 3, 297–312. MR 1635264 4, 131

[58] János Komlós, Ali Shokoufandeh, Miklós Simonovits, and Endre Szemerédi, *The regularity lemma and its applications in graph theory*, Theoretical aspects of computer science (Tehran, 2000), Lecture Notes in Comput. Sci., vol. 2292, Springer, Berlin, 2002, pp. 84–112. MR 1966181 4

[59] Michael Krivelevich, personal communication. 9

[60] ———, *The critical bias for the Hamiltonicity game is* $(1+o(1))n/\ln n$, J. Amer. Math. Soc. **24** (2011), no. 1, 125–131. MR 2726601 9

[61] ———, *Positional games*, Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV, Kyung Moon Sa, Seoul, 2014, pp. 355–379. MR 3727616 9

[62] Michael Krivelevich, Choongbum Lee, and Benny Sudakov, *Resilient pancyclicity of random and pseudorandom graphs*, SIAM J. Discrete Math. **24** (2010), no. 1, 1–16. MR 2600649 21

[63] ———, *Robust Hamiltonicity of Dirac graphs*, Trans. Amer. Math. Soc. **366** (2014), no. 6, 3095–3130. MR 3180741 11

[64] Michael Krivelevich and Benny Sudakov, *Sparse pseudo-random graphs are Hamiltonian*, J. Graph Theory **42** (2003), no. 1, 17–33. MR 1943104 21

[65] Michael Krivelevich, Benny Sudakov, and Tibor Szabó, *Triangle factors in sparse pseudo-random graphs*, Combinatorica **24** (2004), no. 3, 403–426. MR 2085364 5, 7, 21

[66] Harold W. Kuhn, *The Hungarian method for the assignment problem*, Naval Res. Logist. Quart. **2** (1955), 83–97. MR 75510 131

[67] Willem Mantel, *Vraagstuk XXVIII*, Wiskundige Opgaven **10** (1907), 60–61 (Dutch). 120

[68] Richard Montgomery, *Spanning trees in random graphs*, Adv. Math. **356** (2019), 106793, 92. MR 3998769 6

[69] Rajko Nenadov, *Ramsey and universality properties of random graphs*, Ph.D. thesis, ETH Zurich, 2016. 7, 8, 133

[70] Frank P. Ramsey, *On a Problem of Formal Logic*, Proc. London Math. Soc. (2) **30** (1929), no. 4, 264–286. MR 1576401 120

[71] Vojtech Rödl and Andrzej Ruciński, *Perfect matchings in ε-regular graphs and the blow-up lemma*, Combinatorica **19** (1999), no. 3, 437–452. MR 1723256 4, 23

[72] Vojtech Rödl, Andrzej Ruciński, and Anusch Taraz, *Hypergraph packing and graph embedding*, Combin. Probab. Comput. **8** (1999), no. 4, 363–376, Random graphs and combinatorial structures (Oberwolfach, 1997). MR 1723649 4

[73] David Saxton and Andrew Thomason, *Hypergraph containers*, Invent. Math. **201** (2015), no. 3, 925–992. MR 3385638 5

[74] Alexander Scott, *Szemerédi's regularity lemma for matrices and sparse graphs*, Combin. Probab. Comput. **20** (2011), no. 3, 455–466. MR 2784637 4

[75] Benny Sudakov and V. H. Vu, *Local resilience of graphs*, Random Structures Algorithms **33** (2008), no. 4, 409–433. MR 2462249 4, 10, 21

[76] Endre Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 199–245. MR 369312 4

[77] ———, *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Colloq. Internat. CNRS, vol. 260, CNRS, Paris, 1978, pp. 399–401. MR 540024 4

[78] Andrew Thomason, *Pseudorandom graphs*, Random graphs '85 (Poznań, 1985), North-Holland Math. Stud., vol. 144, North-Holland, Amsterdam, 1987, pp. 307–331. MR 930498 6

[79] Paul Turán, *Eine Extremalaufgabe aus der Graphentheorie*, Mat. Fiz. Lapok **48** (1941), 436–452. MR 18405 120

AUTHORS

Peter Allen
Department of Mathematics
London School of Economics
Houghton Street
London WC2A 2AE, U.K.
p.d.allen@lse.ac.uk

Julia Böttcher
Department of Mathematics
London School of Economics
Houghton Street
London WC2A 2AE, U.K.
j.boettcher@lse.ac.uk

Hiệp Hàn
Departamento de Matemática y Ciencia de la Computación
Universidad de Santiago de Chile
Las Sophoras 173
Estación Central, Santiago, Chile
hiep.han@usach.cl

Yoshiharu Kohayakawa
Instituto de Matemática e Estatística
Universidade de São Paulo
Rua do Matão 1010
05508–090 São Paulo, Brazil
yoshi@ime.usp.br

Yury Person
Institut für Mathematik
Technische Universität Ilmenau
98684 Ilmenau, Germany
yury.person@tu-ilmenau.de