

PAPER

Finding independent transversals efficiently

Alessandra Graf^{*} and Penny Haxell¹

Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON, N2L 3G1, Canada

^{*}Corresponding author. Email: agraf@uwaterloo.ca

(Received 6 November 2018; revised 26 February 2020; accepted 6 March 2020; first published online 14 May 2020)

Abstract

We give an efficient algorithm that, given a graph G and a partition V_1, \dots, V_m of its vertex set, finds either an *independent transversal* (an independent set $\{v_1, \dots, v_m\}$ in G such that $v_i \in V_i$ for each i), or a subset \mathcal{B} of vertex classes such that the subgraph of G induced by $\bigcup \mathcal{B}$ has a small dominating set. A non-algorithmic proof of this result has been known for a number of years and has been used to solve many other problems. Thus we are able to give algorithmic versions of many of these applications, a few of which we describe explicitly here.

2020 MSC Codes: Primary 05C69; Secondary 05C85

1. Introduction

Let G be a graph whose vertex set is partitioned into classes V_1, \dots, V_m . An *independent transversal* (IT) of G with respect to the given vertex partition is an independent set $\{v_1, \dots, v_m\}$ in G such that $v_i \in V_i$ for each i . This is a very general notion, and many combinatorial problems can be formulated by asking if a given graph with a given vertex partition has an IT. Indeed, the SAT problem can be formulated in these terms (see *e.g.* [41]), and so we cannot expect to find an efficient characterization of those G for which an IT exists. However, there are now various known results giving sufficient conditions for the existence of an IT. One of the most easily stated and most frequently applied is the following result from [36, 38].

Theorem 1.1. *Let G be a graph with maximum degree Δ . Then, for any vertex partition (V_1, \dots, V_m) of G where $|V_i| \geq 2\Delta$ for each i , there exists an independent transversal of G .*

Theorem 1.1 answered the question of how large the vertex classes need to be, in terms of the maximum degree, to guarantee the existence of an IT in G . This question was first introduced and studied in 1975 by Bollobás, Erdős and Szemerédi [21], and further progress was contributed over the years by many authors. In particular, linear upper bounds in terms of Δ were given by Alon [7] (in an early application of the Lovász Local Lemma [26]) and independently Fellows [27], and a later application of the Local Lemma gave that class size $2e\Delta$ is sufficient (see *e.g.* Alon and Spencer [13]). Further refining this approach, Bissacot, Fernández, Procacci and Scoppola [20] improved this to 4Δ . Work on lower bounds included results of Jin [44], Yuster [60] and Alon [10], and in 2006 Szabó and Tardos [59] gave constructions for every Δ in which $|V_i| = 2\Delta - 1$

¹Partially supported by NSERC.




Figure 1. A constellation K for the set \mathcal{B} of classes enclosed by the dotted border. Each circle represents a vertex class. The centres of the stars appear in black (and have labels) and the leaves appear in red.

for each i but there is no IT. Therefore Theorem 1.1 is best possible for every value of Δ . A more precise version of Theorem 1.1 (also involving the number m of vertex classes) is given in [39].

Theorem 1.1 is an immediate consequence of a more general statement described in terms of domination (stated explicitly in [38], although it follows easily from the argument in [36]). We say that a subset $D \subseteq V(G)$ dominates a subgraph W of G if, for all $w \in V(W)$, there exists $uw \in E(G)$ for some $u \in D$. (This definition of domination is quite often referred to as *strong domination* or *total domination*, but since it is the only notion of domination that we will refer to in this paper, we will use the simpler term.) For a vertex partition (V_1, \dots, V_m) of G and a subset \mathcal{B} of $\{V_1, \dots, V_m\}$, we write $G_{\mathcal{B}}$ for the subgraph $G[\bigcup_{V_i \in \mathcal{B}} V_i] - \{uv \in E(G) : u, v \in V_i \text{ for some } V_i \in \mathcal{B}\}$ obtained by removing any edges inside vertex classes from the subgraph of G induced by the union of the classes in \mathcal{B} .

Theorem 1.2. *Let G be a graph with a vertex partition (V_1, \dots, V_m) . Suppose that, for each $\mathcal{B} \subseteq \{V_1, \dots, V_m\}$, the subgraph $G_{\mathcal{B}}$ is not dominated in $G_{\mathcal{B}}$ by any set of size at most $2(|\mathcal{B}| - 1)$. Then G has an IT.*

To see that Theorem 1.2 implies Theorem 1.1, simply note that if the union of $|\mathcal{B}|$ vertex classes in G contains a total of $2\Delta|\mathcal{B}|$ vertices then $G_{\mathcal{B}}$ cannot be dominated by $2|\mathcal{B}| - 2$ vertices of degree at most Δ .

In fact the proof of Theorem 1.2 shows that if G does not have an IT then there exists $\mathcal{B} \subseteq \{V_1, \dots, V_m\}$ such that $G_{\mathcal{B}}$ is dominated by the vertex set of a constellation for \mathcal{B} .

Definition 1.3. *Let \mathcal{B} be a set of vertex classes in a vertex-partitioned graph G . A constellation for \mathcal{B} is an induced subgraph K of $G_{\mathcal{B}}$, whose components are stars with at least two vertices, each with a centre and a non-empty set of leaves distinct from its centre. The set of all leaves of K forms an IT of $|\mathcal{B}| - 1$ vertex classes of \mathcal{B} .*

Note that if K is a constellation for \mathcal{B} then $|V(K)| \leq 2(|\mathcal{B}| - 1)$. Figure 1 shows an example of a constellation.

Theorem 1.2 (often in the form of Theorem 1.1) has been used to obtain many results in various fields, including graph theory (e.g. list colouring [38], strong colouring [4, 39], delay edge colouring [11], circular colouring [45, 46], various graph partitioning and special independent set problems [12, 24, 43, 47]), hypergraphs (e.g. hypergraph matching [14, 15, 36, 50]), group theory (e.g. generators in linear groups [22]) and theoretical computer science (e.g. job scheduling and other resource allocation problems [17, 18]). Unfortunately the proofs of Theorems 1.1 and 1.2 are not algorithmic. For certain applications it is enough to know that class sizes of $c\Delta$ in Theorem 1.1 guarantee an IT for some constant c , so for these one could obtain algorithmic versions using (for

example) the algorithmic Lovász Local Lemma [19, 53, 54]. However, for many other applications, having the best possible value of the constant c is important. This therefore raises the question of how much the hypotheses in Theorems 1.2 and 1.1 need to be strengthened in order to guarantee that an IT can be found efficiently, thereby giving algorithmic proofs of these applications with the constants being as close as possible to their optimal values.

Most of the known results on this question have focused on Theorem 1.1, and have been obtained as applications of algorithmic versions of the Lovász Local Lemma or its lopsided variant. These include the original algorithm of Beck [19] and its improvements (see e.g. [8, 20, 25, 52, 58]) as well as the resampling algorithm of Moser and Tardos [53] and its improvements (see e.g. [1, 33, 34, 48, 49, 55]). In particular, using the Moser–Tardos approach (and based on [20, 55]), Harris and Srinivasan [35] gave a randomized algorithm that finds an IT in expected time $O(m\Delta)$ in graphs with class size 4Δ . The current best result for polynomial expected time is due to Harris [31] who improved the bound on the class size to $4\Delta - 1$. Deterministic algorithms based on derandomizing the Moser–Tardos algorithm have also been studied, but they require the class size to be $C\Delta$ for some large constant C in order to find an IT efficiently [28, 32]. Some of these deterministic algorithms are known to be parallelizable [23, 32].

In this paper we address the algorithmic IT question for a large class of graphs, without using the Lovász Local Lemma or any of its variants. A graph G with vertex partition (V_1, \dots, V_m) is said to be r -claw-free with respect to (V_1, \dots, V_m) if no vertex of G has r independent neighbours in distinct vertex classes. Our main theorem is as follows.

Theorem 1.4. *There exists an algorithm FindITorBD that takes as input $r \in \mathbb{N}$ and $\varepsilon > 0$, and a graph G with vertex partition (V_1, \dots, V_m) such that G is r -claw-free with respect to (V_1, \dots, V_m) , and finds either:*

- (1) *an independent transversal in G , or*
- (2) *a set \mathcal{B} of vertex classes and a set D of vertices of G such that D dominates $G_{\mathcal{B}}$ in G and $|D| < (2 + \varepsilon)(|\mathcal{B}| - 1)$. Moreover, D contains $V(K)$ for a constellation K for some $\mathcal{B}_0 \supseteq \mathcal{B}$, where $|D \setminus V(K)| < \varepsilon(|\mathcal{B}| - 1)$.*

For fixed r and ε , the running time is polynomial in $|V(G)|$.

Note in particular that any graph with maximum degree Δ is $(\Delta + 1)$ -claw-free with respect to any partition. Thus, taking $r = \Delta + 1$ and $\varepsilon = 1/\Delta$ gives the following algorithmic version of Theorem 1.1.

Corollary 1.5. *Let $\Delta \in \mathbb{N}$ be given. Then there exists an algorithm that takes as input any graph G with maximum degree Δ and vertex partition (V_1, \dots, V_m) such that $|V_i| \geq 2\Delta + 1$ for each i and finds, in time polynomial in $|V(G)|$, an independent transversal in G .*

Therefore only a slight strengthening of the hypotheses is required to make these results algorithmic.

As shown in the proof of Theorem 1.4 in Section 6, the running time of the algorithm FindITorBD is $O(|V(G)|^g)$, where g is a function of r and ε . Similarly, for Corollary 1.5 the degree g depends on Δ . We remark that, for simplicity, in this paper we make no attempt to optimize the running time of our algorithms in terms of these parameters.

The proof of Theorem 1.4 explicitly describes the algorithm FindITorBD. It uses ideas from the original (non-algorithmic) proof of Theorem 1.2 (see [36, 43]), and modifications of several key notions (including that of ‘lazy updates’) introduced by Annamalai [14, 16], who gave an algorithmic version of the specific case of matchings in bipartite hypergraphs. This appears as Theorem 7.2 in Section 7, and is relevant to other well-studied problems such as the restricted max-min fair allocation problem (also known as the Santa Claus problem); see [17, 18]. Theorem 1.4 is a broad

generalization of Theorem 7.2 which, because of the large number of applications of Theorems 1.1 and 1.2 over the years, has algorithmic consequences for many results in a wide variety of settings. In addition to describing the case of bipartite hypergraph matching, in Section 7 we outline algorithmic versions of a few more selected applications of Theorem 1.2 that follow from our results. Here we have chosen to discuss circular chromatic index (Kaiser, Král and Škrekovski [45]), strong colouring (Aharoni, Berger and Ziv [4]), and hitting sets for maximum cliques (King [47]), but there are many other examples, some of which are described in detail in [29]. In each case, the algorithmic version is only slightly weaker than the original result due to the error ε introduced in Theorem 1.4. In fact, for some applications (for example the results on circular chromatic index and hitting sets for maximum cliques) no weakening at all is needed.

This paper is organized as follows. In Section 2 we give an overview of the proof of our main result, by first outlining the proof of Theorem 1.2 (which gives an exponential algorithm) and then sketching how we modify it to make the algorithm efficient. Our algorithms are described in detail in Section 4, after definitions and other preliminary material in Section 3. The main components of the algorithm are analysed in Section 5 and the running time in Section 6. Section 7 is devoted to applications of our results and Section 8 contains concluding remarks and open questions.

2. Set-up and overview

Throughout this paper we will work with the following notation and assumptions. Let r and ε be fixed, let G be a graph and let (V_1, \dots, V_m) be a vertex partition of G such that G is r -claw-free with respect to (V_1, \dots, V_m) . By deleting the edges between vertices in the same vertex class V_i and considering the resulting graph G' , we may assume without loss of generality that each vertex class V_i is an independent set of vertices. This is because a set M is an IT of G' if and only if M is an IT of G . Since the case $m = 1$ is trivial we may assume from now on that $m \geq 2$.

Our algorithms will seek to construct an IT of G step by step, by augmenting a previously constructed *partial independent transversal* (PIT) of G with respect to the given vertex partition. A PIT is simply an independent set M in G (of size at most m) such that no two vertices of M are in the same vertex class.

Note that any isolated vertex can be added to any PIT that does not contain a vertex in its vertex class. Thus we may remove the vertex classes from V_1, \dots, V_m that contain at least one isolated vertex and consider the induced subgraph of the remaining vertex classes as G under the same partition of these vertices. We will therefore assume from now on that G does not contain an isolated vertex. In particular, we may also assume that $r \geq 2$.

We denote the vertex class that contains the vertex $v \in V(G)$ by $A(v)$ and the set of vertex classes containing $W \subseteq V(G)$ by $A(W) = \{A(v) : v \in W\}$. We write $N(v)$ for the neighbourhood in G of v , and $N_W(v)$ for $N(v) \cap W$. We denote $|N_W(v)|$ by $d_W(v)$.

To give an overview of the proof of Theorem 1.4, we first sketch the proof of Theorem 1.2. The proof does give a procedure for constructing an IT, but (as we will note after the sketch) the number of steps could be as large as $(r - 1)^m$.

Sketch of the proof of Theorem 1.2. Let M be a PIT and let A be a vertex class such that $A \cap M = \emptyset$. We aim to alter M until it can be augmented by a vertex of A .

We build a ‘tree-like structure’ T (which we describe as a vertex set inducing a forest of stars) as follows. Choose $x_1 \in A$ and set $T = \{x_1\}$. If $d_M(x_1) = 0$, then *improve* M by adding x_1 to M , and stop. Otherwise add $N_M(x_1)$ to T .

In the general i th step, it can easily be shown that $|T| \leq 2(|A(T)| - 1)$. Thus, by assumption, the subgraph $G_{A(T)}$ of G induced by $\bigcup_{V_i \in A(T)} V_i$ is not dominated by T . Therefore there exists a vertex of $G_{A(T)}$ that is not adjacent to any vertex in T . Choose such a vertex x_i arbitrarily.

If $d_M(x_i) = 0$, then *improve* M by adding x_i to M and removing the M -vertex y in $A(x_i)$ (if it exists). This forms a new PIT M , and is an improvement in the following sense: it reduces $d_M(x_i)$

where y was added to T because it was in $N_M(x_j)$ in an earlier step. Truncate T to $\{x_1\} \cup N_M(x_1) \cup \dots \cup \{x_j\} \cup N_M(x_j)$.

Otherwise add x_i and $N_M(x_i)$ to T . Thus $|A(T)|$ increases by $d_M(x_i) > 0$ and $|T|$ increases by $d_M(x_i) + 1$, thus maintaining $|T| \leq 2(|A(T)| - 1)$. See Figure 1 for T (the set of vertices shown) and $A(T)$ (the set of classes enclosed by the dotted border). Note that T is a constellation for $A(T)$, whose centres are the x_i and whose leaves are the $N_M(x_i)$.

At each step we either grow T OR reduce $d_M(x_j)$ for some j , UNTIL the current M can be extended to include a vertex of A . Thus progress can be measured by a *signature vector*

$$(d_M(x_1), \dots, d_M(x_t), \infty)$$

that has length at most $|M| + 1 \leq m$ since each $N_M(x_i)$ is a non-empty subset of M and all such sets are mutually disjoint. Each step reduces the lexicographic order of the signature vector. Thus the process terminates, and we succeed in extending M to a larger PIT and eventually to an IT. □

The drawback of the above procedure is that the number of signature vectors (and hence the number of steps) could potentially be as large as $(r - 1)^m$, where G is r -claw-free. To make this approach into an efficient algorithm, we make three main modifications. Here the idea of ‘lazy updates’ from [14, 16] is used, which essentially amounts to performing updates in ‘clusters’ (large subsets of vertices) rather than at the level of individual vertices (that change the quantities $d_M(x_i)$ only one at a time). These modifications are as follows.

- (1) *Maintaining layers.* At each growth step, instead of choosing x_i arbitrarily, choose it to be a vertex in a class at smallest possible ‘distance’ from the root class A , similar to a breadth-first search. Vertices x_i added into classes at the same distance from A are in the same *layer* (see Figure 2).
- (2) *Updating in ‘clusters’.* Instead of updating M when a single x_i satisfies $d_M(x_i) = 0$, update only when at least a *positive proportion* μ of an entire layer satisfies $d_M(x) = 0$. Discard later layers.
- (3) *Rebuilding layers in ‘clusters’.* After an update, add new vertices x_i to a layer of T only if doing so would add a μ proportion of that layer. Then discard later layers.

The parameter μ is a fixed positive constant, chosen to be small enough with respect to the parameters ε and $1/r$. The extra ε factor in (2) of Theorem 1.4 is enough to guarantee that the same proof idea as for Theorem 1.2 finds an IT in G , UNLESS (as in that proof) at some point in the construction of T (a subset \mathcal{B} that is almost all of), $A(T)$ is dominated by T (plus a certain very small set of additional vertices, necessary to deal with the error introduced by the lazy updates in modifications 2 and 3). This dominating set will have total size less than $(2 + \varepsilon)(|\mathcal{B}| - 1)$, resulting in output (2) in Theorem 1.4. Thus for the rest of this section we will assume that this never occurs, and in particular that a positive proportion (depending on ε) of the vertices of $G_{A(T)}$ are not dominated by T .

A consequence of *maintaining layers* is that the vertices in $G_{A(T)}$ that do not have a neighbour in T tend to be ‘pushed’ towards the bottom layer. This results in the set of vertices of T in the bottom layer having size a positive proportion ρ of $|T|$, where again ρ depends only on ε and r (see Lemma 5.3). This implies that the total number of layers is always logarithmic in m (Lemma 5.4), since with each new layer the total size of T increases by a fixed factor larger than one.

Updating in clusters and *rebuilding layers in clusters* allow a different signature vector, that measures sizes of layers rather than degrees of individual vertices $d_M(x_i)$. It has just two entries per layer: the first is essentially $-\lceil \log x \rceil$, where x is the number of vertices x_i in that layer, and the second is essentially $\lfloor \log y \rfloor$, where y is the total size of their neighbourhoods in M . Updating M in a cluster (modification 2) decreases the value of y for a layer by a positive proportion. Rebuilding a layer in a cluster (modification 3) increases the value of x for a layer by a positive proportion.

Hence (with suitably chosen bases for the logarithms) these updates always *decrease* the relevant entry by an integer amount. Therefore, as in the proof of Theorem 1.2, each update decreases the signature vector lexicographically (Lemma 6.2).

Since the length of the signature vector is proportional to the number of layers, as noted above this is logarithmic in m . The entries are also of the order $\log m$. While this gives a very significant improvement over the signature vector from the proof of Theorem 1.2, it still does not quite give a polynomial number of signature vectors. However, as in [14, 16] it can be shown that, with a suitable alteration of the above definition of signature vector (Definition 6.1), each signature vector can be associated with a subset of integers from 1 to x , where x is of order $\log m$. It then follows that the number of signature vectors, and hence the number of iterations of the algorithm, is polynomial in m (Lemma 6.4).

3. Preliminaries

In this section we formalize the main notions we will need. Much of the terminology in this section follows that of [14, 16]. Let G and (V_1, \dots, V_m) be as in Section 2. For the definitions that follow, consider a PIT M of G .

Definition 3.1. A vertex u blocks a vertex v if $u \in M$ and $uv \in E(G)$.

Definition 3.2. A vertex v is immediately addable with respect to M if $v \notin M$ and it has no vertices in $V(G)$ blocking it. For $W \subseteq V(G)$, $I_M(W)$ denotes the set of vertices in W that are immediately addable with respect to M .

Definition 3.3. A layer L of G with respect to a PIT M is a pair (X, Y) where:

- (1) $X \subseteq V(G) \setminus M$,
- (2) X is an independent set,
- (3) $Y \subseteq M$ is the set of blocking vertices of X , and
- (4) every $u \in Y$ is adjacent to exactly one vertex from X .

For an example of a layer, refer to Figure 2. Note that Y is also an independent set since M is an independent set.

Definition 3.4. Let M be a PIT in G and let A be a vertex class in the vertex partition of G that does not contain a vertex in M . An alternating tree T with respect to M and A is a tuple (L_0, \dots, L_ℓ) where $\ell \geq 0$ such that:

- (1) $L_0 = (X_0, Y_0) = (\emptyset, \emptyset)$ and $A(Y_0) := A$,
- (2) $L_i = (X_i, Y_i)$ is a layer for each $1 \leq i \leq \ell$,
- (3) $X_1 \subseteq A$ and $X_i \subseteq \bigcup_{v \in Y_{i-1}} A(v)$ for all $i = 2, \dots, \ell$, and
- (4) $(X_i \cup Y_i) \cap (X_{i'} \cup Y_{i'}) = \emptyset$ for all $i, i' \in \{0, \dots, \ell\}$, $|i - i'| > 0$.

We call A the root of T .

Figure 2 provides an example of an alternating tree.

Let $T = (L_0, \dots, L_\ell)$ be an alternating tree of G with respect to a PIT M and root A . For each $0 \leq j \leq \ell$ we let

$$X_{\leq j} = \bigcup_{i=0}^j X_i.$$




Figure 2. An alternating tree T arising from a graph G and PIT M . The circles are vertex classes, and vertices in the same layer of T are enclosed with a dotted border. Within each layer L_i , the vertices in X_i are shown in black and the vertices in Y_i are shown in red.

Similarly, we define

$$Y_{\leq j} = \bigcup_{i=0}^j Y_i.$$

Note that $A(Y_{\leq \ell})$ is the set of vertex classes intersecting T . It follows from Definitions 3.3 and 3.4 that the subgraph of G induced by $X_{\leq \ell} \cup Y_{\leq \ell} \setminus I_M(X_{\leq \ell})$ is a constellation for $A(Y_{\leq \ell})$ (see Figures 1 and 2).

Our algorithm will make use of fixed constants μ , U and ρ which will be chosen in advance and depend only on the input constants r and ε . The following notion formalizes a suitable choice.

Definition 3.5. Let $r \geq 2$ and $\varepsilon > 0$ be given. We say a tuple (μ, U, ρ) of positive real numbers is feasible for (r, ε) if the following hold:

- (1) $(2 + \varepsilon) \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] > \left(\frac{2 + \mu(r + 2) + \rho(r + 1)}{1 - \mu} \right),$
- (2) $\varepsilon \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] > \frac{\mu(r + 4) + \rho(r + 2)}{1 - \mu},$ and
- (3) $U - \mu\rho > \rho.$

For example,

$$(\mu, U, \rho) = \left(\frac{\varepsilon}{10r}, \frac{10r}{\varepsilon}, \frac{\varepsilon}{10r} \right)$$

is feasible for (r, ε) when $r \geq 2$ and $0 < \varepsilon < 1$. Thus, when r and ε are fixed, the parameters μ , U and ρ may also be taken to be fixed constants. As mentioned in the Introduction, we make no attempt here to choose the constants to optimize the running time (see Lemma 6.4).

The following two definitions will depend on the fixed constants μ and U . Definition 3.6 will apply when X, Y forms a partially built layer $L_{\ell+1}$.

Definition 3.6. Let $T = (L_0, \dots, L_\ell)$ be an alternating tree of G with respect to a PIT M and root A and let $X, Y \subseteq V(G)$. A vertex $v \in A(Y_\ell)$ is an *addable vertex* for X, Y and T if $v \notin Y_\ell \cup X \cup Y$, $|A(v) \cap X| < U$, and there does not exist a vertex $u \in X_{\leq \ell} \cup Y_{\leq \ell} \cup X \cup Y$ such that $uv \in E(G)$.

Definition 3.7. A layer $L_i = (X_i, Y_i)$ is *collapsible* if $I_M(X_i) > \mu|X_i|$.

4. Algorithms

Recall that G is an r -claw-free graph with respect to vertex partition (V_1, \dots, V_m) . Let M be a PIT in G and let A be a vertex class in the vertex partition of G that does not contain a vertex in M . The main idea of the algorithm FindITorBD in Theorem 1.4 is to perform a series of modifications to M that will allow us to augment it with a vertex in A . If we are not successful then we will find a subset of classes (based on an alternating tree of G) that has a small dominating set.

The algorithm FindITorBD is given in Section 6. In the following subsections, we describe three algorithms that are used by FindITorBD. The first two algorithms, called BuildLayer and SuperposedBuild, are used as subroutines in the third algorithm, called GrowTransversal. GrowTransversal appears as the main subroutine of FindITorBD.

Let $T = (L_0, \dots, L_\ell)$ be an alternating tree of G with respect to M and A . Let μ, U and ρ be fixed constants chosen in advance so that (μ, U, ρ) is feasible for (r, ε) (see Definition 3.5).

4.1 BuildLayer

BuildLayer is a subroutine in the main algorithm for augmenting M that helps construct new layers for an alternating tree $T = (L_0, \dots, L_\ell)$. The function takes as inputs T and some $X, Y \subseteq V(G)$ ((X, Y) can be thought of as a ‘partially built’ layer). It then creates a new layer $L_{\ell+1} = (X_{\ell+1}, Y_{\ell+1})$ by augmenting X and Y and returning the resulting pair (X, Y) .

```

1: function BUILD_LAYER( $T, X, Y$ )
2:   while there is a vertex  $v \in A(Y_\ell)$  that is addable for  $X, Y$  and  $T$  do
3:      $X := X \cup \{v\}$ 
4:      $Y := Y \cup \{u \in M : uv \in E(G)\}$ 
5:   end while
6:   return  $(X, Y)$ 
7: end function
    
```

4.2 SuperposedBuild

SuperposedBuild is a subroutine in the main algorithm for augmenting M that, after a modification of M occurs in the algorithm, modifies T so that it remains an alternating tree with respect to the new PIT M . SuperposedBuild possibly augments T by adding some vertices that are no longer blocked due to the modification of M . The function takes as inputs the current T and its number of layers ℓ . It then performs some tests on the layers of T , to see if any X_i could be substantially enlarged, and returns a possibly modified alternating tree to replace T for the next iteration of the main algorithm as well as the number of layers in the returned alternating tree.

```

1: function SUPERPOSEDBUILD( $T, \ell$ )
2:    $i := 1$ 
3:   while  $i \leq \ell$  do
4:      $(X'_i, Y'_i) := \text{BuildLayer}((L_0, \dots, L_{i-1}), X_i, Y_i)$ 
5:     if  $|X'_i| \geq (1 + \mu)|X_i|$  then
6:        $(X_i, Y_i) := (X'_i, Y'_i)$ 
7:        $L_i := (X_i, Y_i)$ 
8:        $T := (L_0, \dots, L_i)$ 
9:        $\ell := i$ 
10:    end if
11:     $i := i + 1$ 
12:  end while
13:  return  $(T, \ell)$ 
14: end function
    
```


4.3 GrowTransversal

GrowTransversal is the main algorithm for augmenting M . It takes as inputs M and A and performs a series of modifications to M until either a vertex in A is added to M or an iteration constructs a layer L_i with too small an X_i relative to the size of T . When GrowTransversal terminates, it returns M , T and a flag variable x as (M, T, x) . The variable x is set to 1 if GrowTransversal terminates due to an iteration constructing a layer L_i with at most $\rho|Y_{\leq i-1}|$ vertices in X_i and is set to 0 if GrowTransversal successfully augments M with a vertex in A . If GrowTransversal returns $(M, T, 1)$, we will show in the next section that T contains a subset \mathcal{B} of vertex classes whose vertices are dominated by a set of fewer than $(2 + \varepsilon)|\mathcal{B}|$ vertices with the properties stated in Theorem 1.4.

```

1: function GROWTRANSVERSAL( $M, A$ )
2:    $L_0 := (X_0, Y_0) := (\emptyset, \emptyset)$ 
3:    $A(Y_0) := A$ 
4:    $L_0 := (X_0, Y_0)$ 
5:    $T := (L_0)$ 
6:    $\ell := 0$ 
7:   while  $A \cap M = \emptyset$  do
8:      $(X_{\ell+1}, Y_{\ell+1}) := \text{BuildLayer}(T, \emptyset, \emptyset)$ 
9:      $L_{\ell+1} := (X_{\ell+1}, Y_{\ell+1})$ 
10:     $T := (L_0, \dots, L_\ell, L_{\ell+1})$ 
11:    if  $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$  then
12:      return  $(M, T, 1)$  and terminate
13:    else
14:       $\ell := \ell + 1$ 
15:      while  $|I_M(X_\ell)| > \mu|X_\ell|$  do
16:        if  $\ell = 1$  then
17:          Augment  $M$  with a vertex from  $I_M(X_1)$ 
18:          return  $(M, T, 0)$  and terminate
19:        else
20:          for all  $w \in Y_{\ell-1}$  such that  $I_M(X_\ell) \cap A(w) \neq \emptyset$  do
21:             $M := (M \setminus \{w\}) \cup \{u\}$  for some arbitrary  $u \in I_M(X_\ell) \cap A(w)$ 
22:             $Y_{\ell-1} := Y_{\ell-1} \setminus \{w\}$ 
23:          end for
24:        end if
25:         $T' := (L_0, \dots, L_{\ell-1})$ 
26:         $\ell' := \ell - 1$ 
27:         $(T, \ell) := \text{SuperposedBuild}(T', \ell')$ 
28:      end while
29:    end if
30:  end while
31: end function

```

The GrowTransversal algorithm begins by initializing the alternating tree T with respect to M and its number of layers ℓ . While A does not contain a vertex in the PIT M , the algorithm repeats a building layer operation (line 8) followed by a loop of collapsing operations (lines 15–28) that modify M when enough immediately addable vertices with respect to M are present in the newly constructed layer. Figure 3 shows an example of one collapse operation (lines 20–23).

Performing one collapse operation can result in more vertices becoming immediately addable with respect to the new M in earlier layers. Hence one collapse operation can lead to a cascade of collapse operations being performed on layers of T . Also, any collapse operation can modify M in such a way that SuperposedBuild can modify a layer (X_i, Y_i) of T to contain a significantly




Figure 3. An example of collapsing layer L_ℓ . In (a), a substantial set $I_M(X_\ell)$ of immediately addable vertices is found in X_ℓ (line 15). In (b), the vertex w of M in each class that intersects $I_M(X_\ell)$ is replaced by some $u \in I_M(X_\ell)$ in that same class (line 21). In (c), each $w \in M$ that was replaced in line 21 disappears from $Y_{\ell-1}$ (line 22) and the entire last layer (X_ℓ, Y_ℓ) is removed (line 25).

larger number of vertices ($|X'_i| \geq (1 + \mu)|X_i|$ and so replacing (X_i, Y_i) by (X'_i, Y'_i) adds at least $\mu|X_i|$ vertices to the layer). These modifications to a layer (which is the final layer of the tree returned by SuperposedBuild) may still leave it collapsible, so more collapse operations may be possible if SuperposedBuild modifies a layer of T . Thus, for the T at the conclusion of line 28 to remain an alternating tree with respect to M after all these modifications are made to M , GrowTransversal removes all layers of T constructed after the earliest layer that contains a vertex class where M is modified (lines 25–27). This leaves the resulting T as an alternating tree with respect to the modified M .

Due to these operations, at the beginning and end of each iteration of the main loop of GrowTransversal (which starts in line 7), M remains a PIT and T remains an alternating tree with respect to M . Also, the modifications to M do not change the set of vertex classes containing vertices in M (lines 20–23) unless L_1 is collapsible. In this case, a vertex in $I_M(X_1)$, which is therefore in A , is added to M (lines 15–18) and so A is added to the set of vertex classes covered by the PIT M . As this is the goal of GrowTransversal, the algorithm returns $(M, T, 0)$ and terminates (see line 18).

The algorithm also terminates if, at the start of some iteration, BuildLayer produces a layer $L_{\ell+1}$ whose $X_{\ell+1}$ is not sufficiently large with respect to the number of vertex classes already in T , i.e. $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$ (see lines 11–12). If GrowTransversal terminates because of this, GrowTransversal returns $(M, T, 1)$, which distinguishes this case from when the algorithm terminates because it successfully augments M to include a vertex of A . We will show in the next section that, given an alternating tree $T = (L_0, \dots, L_{\ell+1})$ with respect to M and A such that $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$, there exists some set \mathcal{B} of the vertex classes in T such that $G_{\mathcal{B}}$ is dominated by a set D of vertices with the properties stated in outcome (2) of Theorem 1.4. Our analysis will provide a specific \mathcal{B} and its corresponding D given that GrowTransversal returns $(M, T, 1)$.

5. Analysis

The main result of this section (Lemma 5.3) shows that if GrowTransversal terminates because BuildLayer constructs a layer $L_{\ell+1}$ with $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$ (line 11), then we can find a set \mathcal{B} of vertex classes and a set D of vertices that satisfy the conditions of outcome (2) of Theorem 1.4. Otherwise GrowTransversal will succeed in augmenting the current PIT M with a vertex of A .

As discussed in Section 2, we will bound the total number of steps taken by GrowTransversal using a signature vector (see Definition 6.1) that is defined in terms of logarithmic functions of the sizes of the layers of T . The last lemma of this section (Lemma 5.4) establishes that the number of layers in T (and hence length of any signature vector) is bounded by a logarithmic function of the number m of vertex classes. This fact will be key to the proof given in Section 6 that the total number of signature vectors (and hence the number of steps taken by GrowTransversal) is polynomial in m .

Recall from Section 3 that we can assume G is an r -claw-free graph with respect to vertex partition (V_1, \dots, V_m) , each vertex class V_i is an independent set of vertices, and $r \geq 2$ and $\varepsilon > 0$ are fixed. Also, recall from Section 4 that μ, U and ρ are fixed constants such that (μ, U, ρ) is feasible for (r, ε) .

For concreteness, we may also assume that the vertices of the graph G have been assigned some arbitrary but fixed ordering, and that vertices are processed by our algorithms subject to this ordering. For example, we assume that $\text{BuildLayer}(T, X, Y)$ adds vertices to X and Y in order, that is, during an iteration of the while loop in BuildLayer , the addable vertex with the lowest index in the ordering is the vertex chosen to be v in line 3. Similarly we may assume that the vertex u in line 21 of GrowTransversal is chosen to be the vertex in $I_M(X_\ell) \cap A(w)$ with the lowest index in the ordering. The proof of Lemma 5.3 will also use this convention.

We begin by establishing two preliminary results on basic properties of the alternating tree constructed in GrowTransversal .

Lemma 5.1. *Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree and partial independent transversal at the beginning of some iteration of the while loop in line 7 of GrowTransversal . Then none of the layers L_0, \dots, L_ℓ are collapsible. Hence $|Y_i| \geq (1 - \mu)|X_i|$ for each $i \in \{1, \dots, \ell\}$.*

Proof. Suppose the statement holds at the beginning of the current iteration. In lines 8 and 9 of GrowTransversal , $L_{\ell+1}$ is constructed. If GrowTransversal does not terminate and $L_{\ell+1}$ is not collapsible, then the claim follows for the beginning of the next iteration since none of the earlier layers were modified in the current iteration.

If GrowTransversal does not terminate and $L_{\ell+1}$ is collapsible, then let $T := (L_0, \dots, L_k)$ be the alternating tree with respect to M that results from the loop of collapsing operations where $k \leq \ell$ (i.e. T is the alternating tree after lines 15–28 are completed). Unless the algorithm terminates, $k \geq 1$ (lines 15–18).

Each time a collapse operation is performed (lines 20–23), the number of layers in T is reduced (line 25). Also, if SuperposedBuild modifies T in line 27, the number of layers of T is reduced or stays the same. Thus layers L_0, \dots, L_{k-1} are not changed by the loop of collapsing operations and so remain unchanged throughout the current iteration. Layer L_k may be modified in lines 21, 22 and 27. However, L_k cannot be collapsible since it is the final layer in T after the loop of collapsing operations terminates. Hence none of the layers in an alternating tree are collapsible at the end of an iteration of the while loop in line 7 (unless the algorithm terminates during the iteration).

Since the claim holds for the first iteration of the while loop in line 7, the statement follows by induction on the number of iterations of this loop in GrowTransversal . As L_i is a layer for each $i \in \{1, \dots, \ell\}$, Y_i contains all blocking vertices of all vertices in X_i and every vertex in Y_i is adjacent to exactly one vertex in X_i by construction. Thus there are at most $|Y_i|$ vertices in $X_i \setminus I_M(X_i)$ and at most $\mu|X_i|$ vertices in $I_M(X_i)$ for each layer $i = 1, \dots, \ell$. Hence $|X_i| \leq |Y_i| + \mu|X_i|$ for each $i \in \{1, \dots, \ell\}$. □

Lemma 5.2. *Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree and partial independent transversal at the beginning of some iteration of the while loop in line 7 of GrowTransversal . Then, for each $i \in \{1, \dots, \ell\}$,*

$$(X'_i, Y'_i) := \text{BuildLayer}((L_0, \dots, L_{i-1}), X_i, Y_i)$$

satisfies $|X'_i| < (1 + \mu)|X_i|$.

Proof. Consider layer L_i at the beginning of the current iteration for some $1 \leq i \leq \ell$. During the iteration that L_i was constructed by either BuildLayer in line 8 or SuperposedBuild in line 27, an (additional) application of SuperposedBuild could not increase the size of X_i (both functions created X_i to be as large as possible with respect to M in that iteration).

Suppose no layer built during the iterations between when L_i was constructed and the current iteration is collapsible, that is, the condition of line 15 is not met between the construction of iteration L_i and the current iteration of *GrowTransversal*. Then, since M is not changed during the intervening iterations, *SuperposedBuild* does not increase the size of X_i .

Now suppose some layer built during the intervening iterations was collapsible. Note that the index of the collapsible layer must be greater than i as otherwise L_i would be discarded (lines 25–27). Thus, for each $j \geq i + 1$ such that L_j is collapsible, *SuperposedBuild* tries to augment X_i by at least a μ proportion of its size. However, since L_i is a layer in T at the start of the current iteration, *SuperposedBuild* does not succeed in changing X_i . Hence *BuildLayer*((L_0, \dots, L_{i-1}), X_i, Y_i) does not increase the number of vertices in X_i by $\mu|X_i|$. \square

We are now ready to prove the main result of this section.

Lemma 5.3. *Assume (μ, U, ρ) is feasible for (r, ε) . Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree with root vertex A and partial independent transversal at the beginning of some iteration of the while loop in line 7 of *GrowTransversal*. Then either*

- (1) *when $L_{\ell+1}$ is constructed in line 8 of *GrowTransversal* we have $|X_{\ell+1}| > \rho|Y_{\leq \ell}|$, or*
- (2) **GrowTransversal* terminates in line 12, and there exists a subset \mathcal{B} of the set $\mathcal{B}_0 = A(Y_{\leq \ell})$ such that $G_{\mathcal{B}}$ is dominated by a set D of vertices in G of size less than $(2 + \varepsilon)(|\mathcal{B}| - 1)$. Moreover, $K = G[X_{\leq \ell} \cup Y_{\leq \ell} \setminus I_M(X_{\leq \ell})]$ is a constellation for \mathcal{B}_0 and D contains $V(K)$, where $|D \setminus V(K)| < \varepsilon(|\mathcal{B}| - 1)$.*

Proof. Suppose that after $L_{\ell+1}$ is constructed in line 8, $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$. For each $1 \leq i \leq \ell$, let $(X'_i, Y'_i) = \text{BuildLayer}((L_0, \dots, L_{i-1}), X_i, Y_i)$. Define \mathcal{B} algorithmically by performing the following steps in order:

- (i) $\mathcal{B} := \mathcal{B}_0 = A(Y_{\leq \ell})$,
- (ii) remove all vertex classes of \mathcal{B} that contain U vertices in $X_{\leq \ell+1}$,
- (iii) remove all vertex classes in $\bigcup_{i=1}^{\ell} A(X'_i \setminus X_i)$.

Clearly $\mathcal{B} \subseteq \mathcal{B}_0$, and K is a constellation for \mathcal{B}_0 (as noted after Definition 3.4).

Claim 1. *We have*

$$|\mathcal{B}| - 1 \geq \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] |Y_{\leq \ell}|.$$

Proof. The vertex classes in \mathcal{B} include the vertex classes of T that do not contain U vertices in $X_{\leq \ell+1}$ and do not contain any addable vertices for X_i, Y_i and (L_0, \dots, L_{i-1}) for all $1 \leq i \leq \ell$. We use these facts to bound $|\mathcal{B}|$ from below as follows.

Recall that $|A(Y_{\leq \ell})| = |Y_{\leq \ell}| + 1$. The set A^U of vertex classes in $A(Y_{\leq \ell})$ that contain U vertices in $X_{\leq \ell+1}$ has size at most $U^{-1}|X_{\leq \ell+1}|$. By Lemma 5.2, $|X'_i| < (1 + \mu)|X_i|$ for each $i \in \{1, \dots, \ell\}$. As $X_i \subseteq X'_i$, this implies that there are at most $\mu|X_i|$ vertices in $X'_i \setminus X_i$. Thus $|A(X'_i \setminus X_i)| \leq \mu|X_i|$ for all $1 \leq i \leq \ell$ and so $\sum_{i=1}^{\ell} |A(X'_i \setminus X_i)| \leq \mu|X_{\leq \ell}|$. Also, by Lemma 5.1, $|X_{\leq \ell}| \leq 1/(1 - \mu)|Y_{\leq \ell}|$, and by the assumption, $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$. Therefore

$$\begin{aligned} |\mathcal{B}| &\geq |A(Y_{\leq \ell})| - |A^U| - \left| \bigcup_{i=1}^{\ell} A(X'_i \setminus X_i) \right| \\ &\geq |A(Y_{\leq \ell})| - \frac{1}{U}|X_{\leq \ell+1}| - \mu|X_{\leq \ell}| \end{aligned}$$

$$\begin{aligned}
 &= (|Y_{\leq \ell}| + 1) - \left(\frac{1}{U} |X_{\leq \ell}| + \frac{1}{U} |X_{\ell+1}| + \mu |X_{\leq \ell}| \right) \\
 &\geq |Y_{\leq \ell}| + 1 - \left[\left(\frac{1}{U} + \mu \right) |X_{\leq \ell}| + \frac{\rho}{U} |Y_{\leq \ell}| \right] \\
 &\geq |Y_{\leq \ell}| + 1 - \left[\left(\frac{1}{U} + \mu \right) \left(\frac{1}{1 - \mu} \right) |Y_{\leq \ell}| + \frac{\rho}{U} |Y_{\leq \ell}| \right] \\
 &= 1 + \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] |Y_{\leq \ell}|. \quad \square
 \end{aligned}$$

Let B denote the set of vertices in the vertex classes of \mathcal{B} and let

$$W = X'_{\leq \ell} \cup Y'_{\leq \ell} \cup X_{\ell+1} \cup Y_{\ell+1}.$$

Claim 2. *The set W dominates $G[B \setminus I_M(W)]$.*

Proof. Let $u \in B \setminus I_M(W)$. By (iii) we have that $u \notin X'_{\leq \ell} \setminus X_{\leq \ell}$.

Suppose $u \in W \setminus I_M(W)$. Then $u \in (X_{\leq \ell+1} \cup Y'_{\leq \ell} \cup Y_{\ell+1}) \setminus I_M(W)$. If $u \in Y'_i$ for some $i \in \{1, \dots, \ell\}$, then the construction of (X'_i, Y'_i) and (X_i, Y_i) by BuildLayer implies that u has a neighbour v in X'_i (lines 3–4 of BuildLayer). Hence $v \in W$. Similarly, if $u \in Y_{\ell+1}$, then u has a neighbour v in $X_{\ell+1}$ (lines 3–4 of BuildLayer). If $u \in X_i$ for some $i \in \{1, \dots, \ell + 1\}$, then since $u \notin I_M(W)$, u has a neighbour v that blocks u . By the construction of (X_i, Y_i) (lines 2–4 of BuildLayer), $v \in Y_i$ and so $v \in W$. Therefore every $u \in W \setminus I_M(W)$ has a neighbour in W . Thus we may assume $u \in B \setminus (W \cup I_M(W))$.

Note that each vertex class in \mathcal{B} has at most one vertex in M and these vertices are in $Y'_{\leq \ell}$, hence $u \notin M$. Since $A(u) \in \mathcal{B}$, let L_i be the layer of $T' = (L_0, \dots, L_{\ell+1})$ such that $A(u) \cap Y_{i-1} \neq \emptyset$ for some $1 \leq i \leq \ell + 1$.

Suppose $i < \ell + 1$ and u has no neighbours in $X'_{\leq i} \cup Y'_{\leq i}$. By (ii), $A(u)$ contains fewer than U vertices in X_i . Also, by (iii), $A(u)$ contains no vertices in $X'_i \setminus X_i$. Thus, by Definition 3.6, u is an addable vertex for X_i, Y_i and (L_0, \dots, L_{i-1}) . Hence BuildLayer $((L_0, \dots, L_{i-1}), X_i, Y_i)$ would not stop until either u is added to X'_i or u has a neighbour in $X'_i \cup Y'_i$. As $u \in B \setminus (W \cup I_M(W))$ and $X'_i \subseteq W$, we know that $u \notin X'_i$. Therefore u has a neighbour in $X'_i \cup Y'_i$.

Now suppose $i = \ell + 1$ and u has no neighbours in W . Then, again by (ii) and (iii), $A(u)$ contains fewer than U vertices in X_i and $A(u)$ contains no vertices in $X'_i \setminus X_i$. Thus, by Definition 3.6, u is an addable vertex for $X_{\ell+1}, Y_{\ell+1}$ and (L_0, \dots, L_ℓ) . Hence BuildLayer $((L_0, \dots, L_\ell), \emptyset, \emptyset)$ would not stop until either u is added to $X_{\ell+1}$ or u has a neighbour in $X_{\ell+1} \cup Y_{\ell+1}$. As $u \in B \setminus (W \cup I_M(W))$ and $X_{\ell+1} \subseteq W$, we know that $u \notin X_{\ell+1}$. Therefore u has a neighbour in $X_{\ell+1} \cup Y_{\ell+1}$.

We therefore conclude that W dominates $G[B \setminus I_M(W)]$. □

Define S to be the set of all $u \in V(G)$ for which there exists $v \in I_M(W)$ such that $u \in N(v)$ and u is the neighbour of v with the smallest index in the ordering.

Claim 3. *The set S dominates $I_M(W)$ and*

$$|S| \leq |I_M(W)| < \frac{2\mu + \rho}{1 - \mu} |Y_{\leq \ell}|.$$

Proof. As each $v \in I_M(W)$ has at least one neighbour in G , and the neighbour with the smallest index in the ordering is in S , we have that S dominates $I_M(W)$ and $|S| \leq |I_M(W)|$. By Lemma 5.2, $|X'_{\leq \ell}| < (1 + \mu)|X_{\leq \ell}|$ and so $|X'_{\leq \ell} \setminus X_{\leq \ell}| < \mu|X_{\leq \ell}|$. Also, by definition, $I_M(W) \subseteq X'_{\leq \ell} \cup X_{\ell+1}$

(see Definition 3.2). Thus

$$\begin{aligned}
 |I_M(W)| &= |I_M(X'_{\leq \ell} \cup X_{\ell+1})| \\
 &= |I_M(X'_{\leq \ell})| + |I_M(X_{\ell+1})| \\
 &\leq |I_M(X_{\leq \ell})| + |I_M(X'_{\leq \ell} \setminus X_{\leq \ell})| + |X_{\ell+1}| \\
 &< \mu|X_{\leq \ell}| + \mu|X'_{\leq \ell}| + |X_{\ell+1}| \\
 &= 2\mu|X_{\leq \ell}| + |X_{\ell+1}|.
 \end{aligned}$$

Recall that

$$|X_{\leq \ell}| \leq \frac{1}{1 - \mu}|Y_{\leq \ell}|$$

by Lemma 5.1. Since conclusion (1) of Lemma 5.3 does not hold, we know that $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$, and so we obtain

$$|I_M(W)| < \frac{2\mu}{1 - \mu}|Y_{\leq \ell}| + \rho|Y_{\leq \ell}|,$$

from which the claim follows. □

Let $D = W \cup S$. Then D contains $V(K)$, and by Claim 2 and the choice of S we see that D dominates $G_{\mathcal{B}}$.

To help us estimate the size of D , we first establish the following.

Claim 4. *We have $|Y'_{\leq \ell} \setminus Y_{\leq \ell}| < \mu(r - 1)|X_{\leq \ell}|$.*

Proof. Let $v \in Y'_{\leq \ell} \setminus Y_{\leq \ell}$. Then by definition $v \in M$, and hence v blocks its neighbours. If $uv \in E(G)$ for any $u \in X_i$ such that $1 \leq i \leq \ell$, v blocks u and so would be included in Y_i (line 4 of BuildLayer). This implies $v \in Y_{\leq \ell}$, which is a contradiction. Therefore v is not adjacent to any $u \in X_{\leq \ell}$. However, $v \in Y'_i \setminus Y_i$ for some $1 \leq i \leq \ell$ and, by the construction of (X'_i, Y'_i) , v is adjacent to exactly one $u \in X'_i$ (lines 2–5). Thus v has a neighbour in $X'_i \setminus X_i$ and so v has a neighbour in $X'_{\leq \ell} \setminus X_{\leq \ell}$.

Note that Y'_i is a set of independent vertices in distinct vertex classes. As G is r -claw-free, each vertex of X'_i has at most $r - 1$ independent neighbours in different vertex classes. Since $|X'_{\leq \ell} \setminus X_{\leq \ell}| < \mu|X_{\leq \ell}|$, we have

$$x|Y'_{\leq \ell} \setminus Y_{\leq \ell}| \leq (r - 1)|X'_{\leq \ell} \setminus X_{\leq \ell}| < \mu(r - 1)|X_{\leq \ell}|. \quad \square$$

It remains to bound $|D|$ and $|D \setminus V(K)|$. To do this, we note that $D = X_{\leq \ell} \cup Y_{\leq \ell} \cup Q$ and $D \setminus V(K) = Q \cup I_M(X_{\leq \ell})$ where

$$Q = (X'_{\leq \ell} \setminus X_{\leq \ell}) \cup (Y'_{\leq \ell} \setminus Y_{\leq \ell}) \cup X_{\ell+1} \cup Y_{\ell+1} \cup S.$$

Claim 5. *We have*

$$|Q| < \frac{\mu(r + 2) + \rho(r + 1)}{1 - \mu}|Y_{\leq \ell}|.$$

Proof. Since G is r -claw-free we know that $|Y_{\ell+1}| \leq (r - 1)|X_{\ell+1}|$. Since conclusion (1) of Lemma 5.3 does not hold, we know that $|X_{\ell+1}| \leq \rho|Y_{\leq \ell}|$, and therefore $|Y_{\ell+1}| \leq \rho(r - 1)|Y_{\leq \ell}|$. We bound each of the remaining three summands below using (respectively) Lemma 5.2, Claim 4 and Claim 3, to obtain

$$\begin{aligned}
 |Q| &\leq |X'_{\leq \ell} \setminus X_{\leq \ell}| + |Y'_{\leq \ell} \setminus Y_{\leq \ell}| + |X_{\ell+1}| + |Y_{\ell+1}| + |S| \\
 &\leq \mu|X_{\leq \ell}| + \mu(r-1)|X_{\leq \ell}| + \rho|Y_{\leq \ell}| + \rho(r-1)|Y_{\leq \ell}| + \frac{2\mu + \rho}{1 - \mu}|Y_{\leq \ell}| \\
 &= \mu r|X_{\leq \ell}| + \rho r|Y_{\leq \ell}| + \frac{2\mu + \rho}{1 - \mu}|Y_{\leq \ell}| \\
 &< \mu r|X_{\leq \ell}| + \frac{2\mu + \rho(r+1)}{1 - \mu}|Y_{\leq \ell}|.
 \end{aligned}$$

Since

$$|X_{\leq \ell}| \leq \frac{1}{1 - \mu}|Y_{\leq \ell}|$$

by Lemma 5.1, we conclude

$$|Q| < \frac{\mu(r+2) + \rho(r+1)}{1 - \mu}|Y_{\leq \ell}|$$

as required. □

Now Claim 5 and Lemma 5.1 combine to give

$$\begin{aligned}
 |D| &= |X_{\leq \ell} \cup Y_{\leq \ell} \cup Q| \\
 &< \left(\frac{1}{1 - \mu} + 1 + \frac{\mu(r+2) + \rho(r+1)}{1 - \mu} \right) |Y_{\leq \ell}| \\
 &< \left(\frac{2 + \mu(r+2) + \rho(r+1)}{1 - \mu} \right) |Y_{\leq \ell}|.
 \end{aligned}$$

Claim 1 and the fact that (μ, U, ρ) is feasible for (r, ε) (condition (1) in Definition 3.5) tell us that

$$\begin{aligned}
 (2 + \varepsilon)(|\mathcal{B}| - 1) &\geq (2 + \varepsilon) \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] |Y_{\leq \ell}| \\
 &> \left(\frac{2 + \mu(r+2) + \rho(r+1)}{1 - \mu} \right) |Y_{\leq \ell}| \\
 &> |D|.
 \end{aligned}$$

To bound $|D \setminus V(K)|$ using Claim 5, we observe that $I_M(X_{\leq \ell}) \subseteq I_M(W)$, so we may use Claim 3 to conclude that

$$\begin{aligned}
 |D \setminus V(K)| &= |Q \cup I_M(X_{\leq \ell})| \\
 &< \frac{\mu(r+2) + \rho(r+1)}{1 - \mu}|Y_{\leq \ell}| + \frac{2\mu + \rho}{1 - \mu}|Y_{\leq \ell}| \\
 &= \frac{\mu(r+4) + \rho(r+2)}{1 - \mu}|Y_{\leq \ell}|.
 \end{aligned}$$

Using Claim 1 again, and condition (2) of Definition 3.5 of feasibility of (μ, U, ρ) , we find that

$$\begin{aligned}
 \varepsilon(|\mathcal{B}| - 1) &\geq \varepsilon \left[1 - \frac{1}{U} \left(\frac{1 + \mu U}{1 - \mu} + \rho \right) \right] |Y_{\leq \ell}| \\
 &> \frac{\mu(r+4) + \rho(r+2)}{1 - \mu}|Y_{\leq \ell}| \\
 &> |D \setminus V(K)|.
 \end{aligned}$$

This completes the proof of Lemma 5.3. □

The length of the signature vectors defined in Definition 6.1 of the next section (which will be our measure of the progress of GrowTransversal) depends on the number of layers of the alternating trees T constructed by the algorithm. Our last result of this section gives an upper bound on this quantity.

Lemma 5.4. *Suppose (μ, U, ρ) is feasible for (r, ε) . The number of layers in the alternating tree T with respect to partial independent transversal M maintained during the execution of GrowTransversal is always bounded by $c \log(m)$ where $c = 1/\log[1 + \rho(1 - \mu)]$.*

Proof. Suppose $T = (L_0, \dots, L_\ell)$ at the beginning of some iteration of the while loop in line 7 of GrowTransversal. Consider any L_i where $1 \leq i \leq \ell$. By Lemma 5.1, $|Y_i| > (1 - \mu)|X_i|$. Since GrowTransversal did not terminate in the iteration in which L_i was constructed, by Lemma 5.3 we have

$$|Y_i| > (1 - \mu)|X_i| > \rho(1 - \mu)|Y_{\leq i-1}|.$$

Therefore, since $|Y_1| \geq 1$, we find

$$m \geq |Y_{\leq \ell}| = \sum_{i=0}^{\ell} |Y_i| > [1 + \rho(1 - \mu)]^\ell.$$

Thus the number of layers ℓ at any moment of the algorithm is bounded above by $\log m / \log[1 + \rho(1 - \mu)]$. As μ and ρ are fixed constants, this is $c \log(m)$ for $c = 1/\log[1 + \rho(1 - \mu)]$. □

6. Signatures

Recall from Section 4 that r and ε are fixed constants, and μ, U and ρ are fixed such that (μ, U, ρ) is feasible for (r, ε) . We begin this section by defining the signature vector of an alternating tree. We then use these signature vectors to prove that GrowTransversal terminates after a polynomial in m number of iterations, where the degree of the polynomial is a function of r and ε .

Definition 6.1. *Let $T = (L_0, \dots, L_\ell)$ be an alternating tree with respect to PIT M and vertex class A . The signature of layer L_i is defined to be*

$$(s_{2i-1}, s_{2i}) = \left(- \left\lfloor \log_b \frac{\rho^{-i}}{(1 - \mu)^{i-1}} |X_i| \right\rfloor, \left\lfloor \log_b \frac{\rho^{-i}}{(1 - \mu)^i} |Y_i| \right\rfloor \right)$$

where $b = U/(U - \mu\rho)$. The signature vector of T is $s = (s_1, s_2, \dots, s_{2\ell-1}, s_{2\ell}, \infty)$.

The above definition for the signature of a layer L_i is chosen so that the lexicographic value of the signature vector decreases whenever $|X_i|$ increases significantly (see Lemma 6.2, subcase 2.2) and decreases whenever $|Y_i|$ decreases significantly (Lemma 6.2 subcase 2.1). The factors $\rho^{-i}/((1 - \mu)^{i-1})$ and $\rho^{-i}/((1 - \mu)^i)$ are present to ensure that the coordinates of the signature vector are non-decreasing in absolute value, which we will show in Lemma 6.3. These two properties, together with Lemma 5.4, combine to give the desired upper bound on the total number of signature vectors (Lemma 6.4).

We begin by showing that the lexicographic value of the signature vector decreases during each iteration of GrowTransversal.

Lemma 6.2. *The lexicographic value of the signature vector reduces across each iteration of the loop in line 7 of GrowTransversal unless the algorithm terminates during that iteration.*

Proof. Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree and PIT at the beginning of some iteration of the while loop in line 7 of *GrowTransversal*. Let $s = (s_1, \dots, s_{2\ell}, \infty)$ be the signature vector of T . There are two cases.

Case 1. No collapse operation occurs in this iteration, *i.e.* $|I_M(X_{\ell+1})| \leq \mu|X_{\ell+1}|$.

The only modification to T in this iteration is that a new layer $L_{\ell+1}$ is added to T (lines 8–10). The new signature vector for T is therefore $s' = (s'_1, \dots, s'_{2\ell+2}, \infty)$, where $s'_i = s_i$ for all $1 \leq i \leq 2\ell$ and $(s'_{2\ell+1}, s'_{2\ell+2})$ is the signature of layer $L_{\ell+1}$. Hence the lexicographic value is reduced.

Case 2. At least one collapse operation occurs in this iteration, *i.e.* $|I_M(X_{\ell+1})| > \mu|X_{\ell+1}|$.

Consider the alternating tree T^* returned by *SuperposedBuild* during the last iteration of the loop of collapses (lines 15–28) in this iteration of *GrowTransversal*. Either *SuperposedBuild* returns the same tree as its input T' or *SuperposedBuild* modifies a layer of T and removes all later layers (*i.e.* $T^* = (L_0, \dots, L_k)$ for some $k \leq \ell$ where $L_k = (X'_k, Y'_k)$).

Subcase 2.1. *SuperposedBuild* returns its input T' as T^* .

Let t be the index of the layer of T satisfying $|I_M(X_t)| > \mu|X_t|$ in the final iteration of the loop of collapses. If $t = 1$, a vertex in $I_M(X_1)$ is added to M and the algorithm terminates. Otherwise, $t > 1$ and $T' = (L'_0, \dots, L'_{t-1})$, where $L'_i = L_i$ for all $0 \leq i < t - 1$ and L'_{t-1} is modified from L_{t-1} as lines 20–23 describe.

The signature vector of T' is therefore $s' = (s'_1, \dots, s'_{2t-2}, \infty)$, where $s'_i = s_i$ for all $1 \leq i \leq 2t - 4$ and

$$(s'_{2t-3}, s'_{2t-2}) = \left(- \left\lfloor \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-2}} |X'_{t-1}| \right\rfloor, \left\lfloor \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} |Y'_{t-1}| \right\rfloor \right).$$

Since lines 20–23 do not modify X_ℓ in *GrowTransversal*, X_{t-1} is not modified and so $X'_{t-1} = X_{t-1}$.

Note that each vertex class containing a vertex in $I_M(X_t)$ must contain a vertex in Y_{t-1} . Thus, since a vertex class contains at most U vertices in X_t , there are at least $(\mu/U)|X_t|$ blocking vertices in Y_{t-1} that are not in T' because of the replacements in lines 20–23. By Lemma 5.3 we have $|X_t| > \rho|Y_{\leq t-1}| \geq \rho|Y_{t-1}|$, and so

$$|Y'_{t-1}| \leq |Y_{t-1}| - \frac{\mu}{U}|X_t| < \left(1 - \frac{\mu\rho}{U} \right) |Y_{t-1}|.$$

Therefore

$$\begin{aligned} \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} |Y'_{t-1}| &< \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} \left(1 - \frac{\mu\rho}{U} \right) |Y_{t-1}| \\ &\leq \log_b \left(1 - \frac{\mu\rho}{U} \right) + \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} |Y_{t-1}| \\ &\leq \log_b \left(\frac{U - \mu\rho}{U} \right) + \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} |Y_{t-1}| \\ &= -1 + \log_b \frac{\rho^{-(t-1)}}{(1-\mu)^{t-1}} |Y_{t-1}|. \end{aligned}$$

Hence $s'_{2t-3} = s_{2t-3}$ and $s'_{2t-2} < s_{2t-2}$, so the lexicographic value of the signature vector is reduced.

Subcase 2.2. SuperposedBuild returns an alternating tree T^* different from its input T' .

Let t be the index of the layer of T satisfying $|I_M(X_t)| > \mu|X_t|$ in the final iteration of the loop of collapses. Let q be the index of the last layer in the alternating tree returned by SuperposedBuild in the final iteration of the loop of collapses. Hence $|X'_q| \geq (1 + \mu)|X_q|$.

As (μ, U, ρ) is feasible, $\rho \leq U - \mu\rho$. Hence

$$\frac{U}{U - \mu\rho} = 1 + \frac{\mu\rho}{U - \mu\rho} \leq 1 + \mu.$$

Thus

$$\begin{aligned} \log_b \frac{\rho^{-q}}{(1 - \mu)^{q-1}} |X'_q| &\geq \log_b \frac{\rho^{-q}}{(1 - \mu)^{q-1}} (1 + \mu) |X_q| \\ &= \log_b (1 + \mu) + \log_b \frac{\rho^{-q}}{(1 - \mu)^{q-1}} |X_q| \\ &\geq 1 + \log_b \frac{\rho^{-q}}{(1 - \mu)^{q-1}} |X_q| \end{aligned}$$

and so $s'_{2q-1} < s_{2q-1}$. Since SuperposedBuild and the loop of collapses do not modify layers L_0, \dots, L_{q-1} , we see that $s'_i = s_i$ for all $1 \leq i \leq 2q - 2$. Hence the signature vector of T^* returned by SuperposedBuild is $(s'_1, \dots, s'_{2q-1}, s'_{2q}, \infty)$. Thus the lexicographic value of the signature vector is reduced. □

Lemma 6.3. *The coordinates of the signature vector are non-decreasing in absolute value at the beginning of each iteration of the while loop in line 7 of GrowTransversal.*

Proof. Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree and PIT at the beginning of some iteration of the while loop in line 7 of GrowTransversal. Consider layer L_i for some $1 \leq i \leq \ell$. Since $|Y_i| \geq (1 - \mu)|X_i|$ by Lemma 5.1,

$$|s_{2i-1}| = \left\lfloor \log_b \frac{\rho^{-i}}{(1 - \mu)^{i-1}} |X_i| \right\rfloor \leq \left\lfloor \log_b \frac{\rho^{-i}}{(1 - \mu)^i} |Y_i| \right\rfloor = |s_{2i}|.$$

Hence the coordinates of the signature vector for a layer of T are non-decreasing in absolute value. Now consider layers L_i and L_{i+1} for some $0 \leq i \leq \ell - 1$. Lemma 5.3 implies $|X_i| \geq \rho|Y_{i-1}|$ and so

$$|s_{2i}| = \left\lfloor \log_b \frac{\rho^{-i}}{(1 - \mu)^i} |Y_i| \right\rfloor \leq \left\lfloor \log_b \frac{\rho^{-(i+1)}}{(1 - \mu)^i} |X_{i+1}| \right\rfloor = |s_{2i+1}|.$$

Thus consecutive coordinates of the signature vector for coordinates of different layers of T are also non-decreasing in absolute value. Hence the coordinates of the signature vector of T are non-decreasing in absolute value. □

We may now use Lemmas 6.2 and 6.3 to bound the total number of possible signature vectors.

Lemma 6.4. *Let $T = (L_0, \dots, L_\ell)$ and M be the alternating tree and partial independent transversal at the beginning of some iteration of the while loop in line 7 of GrowTransversal. The number of possible signature vectors for T is bounded by a polynomial in m of degree k where k depends only on r and ε .*

Proof. For each layer of T , the signature vector of T contains two coordinates. Thus, by Lemma 5.4, the signature vector of T has at most $2c \log m$ coordinates where $c = 1/\log [1 + \rho(1 - \mu)]$. Also, by Lemma 6.3, the coordinates are non-decreasing in absolute value

and so the absolute value of the final (finite) coordinate is an upper bound on the absolute value of each coordinate in the signature vector. By Definition 6.1, the final coordinate is

$$\left\lceil \log_b \left(\left[\frac{\rho^{-\ell}}{(1-\mu)^\ell} \right] |Y_\ell| \right) \right\rceil.$$

As $\ell \leq c \log(m)$ (by Lemma 5.4) and $|Y_\ell| \leq m$, the absolute value of each coordinate of the signature vector is bounded above by

$$\begin{aligned} \log_b \left[\frac{\rho^{-c \log(m)}}{(1-\mu)^{c \log(m)}} m \right] &= \log_b m + \log_b \rho^{-c \log(m)} - \log_b (1-\mu)^{c \log(m)} \\ &= \log_b m - c \lceil \log_b(\rho) \rceil \log(m) - c \lceil \log_b(1-\mu) \rceil \log(m) \\ &= \left[\frac{1}{\log(b)} - c \lceil \log_b(\rho) \rceil - c \lceil \log_b(1-\mu) \rceil \right] \log(m). \end{aligned}$$

Let

$$R = \left\lceil \frac{1}{\log(b)} - c \lceil \log_b(\rho) \rceil - c \lceil \log_b(1-\mu) \rceil \right\rceil,$$

and note that R is a fixed constant that depends only on r and ε (since b and c depend only on μ , U and ρ , which in turn depend only on r and ε).

Now to each signature vector $s = (s_1, s_2, \dots, s_{2\ell-1}, s_{2\ell}, \infty)$ we associate the vector $s^+ = (s_1 - 1, s_2 + 2, \dots, s_{2\ell-1} - (2\ell - 1), s_{2\ell} + 2\ell, \infty)$. Then the final coordinate of s^+ is at most $R \log m + 2\ell \leq (R + 2c) \log m$. Since the coordinates of s are non-decreasing in absolute value (and considering the sign pattern), the coordinates of s^+ are strictly increasing in absolute value. Thus each vector s^+ corresponds to a distinct subset of the set $\{1, \dots, \lfloor (R + 2c) \log m \rfloor\}$. Hence the total number of vectors s^+ (and hence the total number of signature vectors) is at most $2^{(R+2c) \log m}$. This completes the proof. \square

(We remark that the idea of the last paragraph was suggested by a referee of [16]; see [15, 16].) We now complete the proof of Theorem 1.4.

Proof of Theorem 1.4. Let FindITorBD be the following algorithm.

```

1: function FINDITORBD( $G; V_1, \dots, V_m$ )
2:    $M := \emptyset$ 
3:   for  $i = 1, \dots, m$  do
4:     Choose an  $A \in \{V_1, \dots, V_m\}$  such that  $A \cap M = \emptyset$ 
5:      $(M, T, x) := \text{GrowTransversal}(M, A)$ 
6:     if  $x = 1$  then
7:        $\mathcal{B} := A(Y_{\leq \ell})$ 
8:       for  $j = 1, \dots, \ell$  do
9:          $(X'_j, Y'_j) := \text{BuildLayer}((L_0, \dots, L_{j-1}), X_j, Y_j)$ 
10:      end for
11:       $\mathcal{B} := \mathcal{B} \setminus A^U$ 
12:       $\mathcal{B} := \mathcal{B} \setminus (\bigcup_{i=1}^\ell A(X'_i \setminus X_i))$ 
13:       $D := X'_{\leq \ell} \cup Y'_{\leq \ell} \cup X_{\ell+1} \cup Y_{\ell+1} \cup S$  as in Lemma 5.3
14:      return  $\mathcal{B}, D$  and ' $G_{\mathcal{B}}$  is dominated by  $D$ ' and terminate
15:    end if
16:     $M := M$  returned by GrowTransversal
17:  end for
18:  return  $M$  and ' $M$  is an independent transversal of  $G$ '
19: end function

```

By Lemma 6.2, every iteration of *GrowTransversal* reduces the lexicographic value of the signature vector of an alternating tree T with respect to a PIT M . Furthermore, Lemma 6.4 implies that the number of such signature vectors is bounded by m^f , where f depends only on r and ε . Thus *GrowTransversal* terminates after at most m^f iterations. It can be easily verified (see [29]) that each iteration of *GrowTransversal* can be implemented to run in time $O(|V(G)|^4)$. Since *GrowTransversal* is implemented at most m times in *FindITorBD*, and the other steps of *FindITorBD* are easily implemented in $O(|V(G)|^4)$ operations, the running time of *FindITorBD* is $O(|V(G)|^4 m^{f+1})$ and thus is polynomial in $|V(G)|$ because r and ε are fixed constants.

It remains to show that *FindITorBD* returns one of the two stated outcomes. *FindITorBD* starts with $M = \emptyset$ and runs *GrowTransversal* at most m times. Note that because of the augmentation in line 17 of *GrowTransversal*, the PIT M at the end of an iteration of *GrowTransversal* covers one more vertex class than the PIT covered at the start of the iteration. Also, the PIT at the end of one iteration of *GrowTransversal* is the initial PIT of the next iteration.

Suppose *FindITorBD* terminates during iteration i of *GrowTransversal*, so that *GrowTransversal* terminates before completing the m th iteration. Then the i th iteration of *GrowTransversal* returned $(M, T, 1)$ for some alternating tree T and PIT M . Since the sets \mathcal{B} and D are defined identically to how they are defined in the proof of Lemma 5.3, we then have that D dominates $G_{\mathcal{B}}$, and (as stated in Lemma 5.3) \mathcal{B} and D have the properties stated in the conclusions of Theorem 1.4.

Suppose *FindITorBD* does not terminate before completing all m iterations of *GrowTransversal*. Then all m vertex classes contain a vertex in the final ‘partial’ IT M . Hence M is an IT of G . Therefore *FindITorBD* returns an IT in G . \square

7. Applications

In this section we briefly discuss some applications of Theorem 1.4. In particular, we discuss applications to hypergraph matching (Section 7.1), circular chromatic index (Section 7.2), strong colouring (Section 7.3) and hitting sets for maximum cliques (Section 7.4). Precise details of these applications, as well as others, appear in [29].

7.1 Hypergraph matchings

Here we consider a hypergraph version of Hall’s theorem for bipartite graphs. An r -uniform bipartite hypergraph $H = (A, B, E)$ is a hypergraph on a vertex set that is partitioned into two sets A and B such that $|e \cap A| = 1$ and $|e \cap B| = r - 1$ for each edge $e \in E$. A perfect matching in H is a subset $M \subseteq E$ of pairwise disjoint edges of H that saturates A , in other words $|M| = |A|$. For a set $S \subseteq A$, we write $E_S = \{e \in E : |e \cap S| = 1\}$ for the set of hyperedges in H incident to S . For a collection of edges $F \subseteq E$, we let $\tau_B(F)$ denote the smallest cardinality of a B -cover of F , that is, a subset $T \subseteq B$ such that $|e \cap T| \neq \emptyset$ for each $e \in F$. The following generalization of Hall’s theorem from [36] provides a condition under which H admits a perfect matching.

Theorem 7.1. *Let $H = (A, B, E)$ be an r -uniform bipartite hypergraph. If*

$$\tau_B(E_S) > (2r - 3)(|S| - 1)$$

for all $S \subseteq A$, then H admits a perfect matching.

It was shown in [36] that Theorem 7.1 is best possible for every r . Note that when $r = 2$ it is (the non-trivial direction of) Hall’s theorem.

In fact Theorem 7.1 is a special case of Theorem 1.2, as can be seen from the following argument. Given $H = (A, B, E)$, construct an auxiliary graph G^H with vertex set E , in which vertices e

and f are adjacent if and only if $e \cap f \cap B \neq \emptyset$. Consider the vertex partition of G^H given by assigning e and f to the same vertex class if and only if $e \cap f \cap A \neq \emptyset$. Thus the vertex classes are indexed by A . With these definitions, a set $M \subseteq E$ is a perfect matching of H if and only if $M \subseteq V(G^H)$ is an IT of G^H .

By Theorem 1.2 applied to G^H (using also the comment before Definition 1.3), if H does not have a perfect matching, then there exists a subset \mathcal{B} of vertex classes (indexed by a set $S(\mathcal{B}) \subseteq A$) that is dominated by $V(K)$ for a constellation K of \mathcal{B} . Thus $T = \bigcup_{e \in V(K)} e \cap B$ is a set of vertices of H forming a B -cover of $E_{S(\mathcal{B})}$. Then the following claim gives an immediate contradiction to the assumption of Theorem 7.1, thus completing the proof.

Claim 6. *Let \mathcal{B} be a subset of the vertex classes of G^H , and let K be constellation for \mathcal{B} . Then*

$$\left| \bigcup_{e \in V(K)} e \cap B \right| \leq (2r - 3)(|S(\mathcal{B})| - 1).$$

Proof. Each component C of K corresponds to a set of edges of H , consisting of the centre e_C of the star C and a non-empty set L_C of leaves, all of which intersect e_C in B . Hence the total number of vertices of B contained in $\{e_C\} \cup L_C$ is at most $(r - 1) + (r - 2)|L_C|$. By definition $\bigcup_C L_C$ is an IT of $|\mathcal{B}| - 1$ classes of \mathcal{B} , implying that K has at most $|S(\mathcal{B})| - 1$ components, and that $\sum_C |L_C| = |S(\mathcal{B})| - 1$. Therefore

$$\begin{aligned} \left| \bigcup_{e \in V(K)} e \cap B \right| &\leq \sum_C (r - 1 + (r - 2)|L_C|) \\ &\leq (r - 1)(|S(\mathcal{B})| - 1) + (r - 2)(|S(\mathcal{B})| - 1) \\ &= (2r - 3)(|S(\mathcal{B})| - 1), \end{aligned}$$

where the sum is over all components C of K . □

The following algorithmic version of Theorem 7.1 was proved by Annamalai in [14, 16].

Theorem 7.2. *For every fixed choice of $r \geq 2$ and $\varepsilon > 0$, there exists an algorithm \mathcal{A} that finds, in time polynomial in the size of the input, a perfect matching in r -uniform bipartite hypergraphs $H = (A, B, E)$ satisfying*

$$\tau_B(E_S) > (2r - 3 + \varepsilon)(|S| - 1)$$

for all $S \subseteq A$.

Here we show that Theorem 1.4 is a generalization of Theorem 7.2. First note that for every r -uniform bipartite hypergraph $H = (A, B, E)$, the graph G^H is r -claw-free with respect to any partition. Indeed, the neighbours of e forming any independent set in G^H must all contain distinct vertices of $e \cap B$, and $|e \cap B| = r - 1$. Thus, given r and ε , we may apply Theorem 1.4 with r and $\varepsilon' = \varepsilon / (r - 1)$ to obtain a polynomial-time algorithm \mathcal{A} that finds for each input $H = (A, B, E)$ either

- (1) an IT in G^H , which is a perfect matching in H , or
- (2) a set \mathcal{B} of vertex classes and a set D of vertices of G^H such that D dominates $G_{\mathcal{B}}$ in G^H and $|D| < (2 + \varepsilon')(|\mathcal{B}| - 1)$. Moreover, D contains $V(K)$ for a constellation K of some $\mathcal{B}_0 \supseteq \mathcal{B}$, where $|D \setminus V(K)| < \varepsilon'(|\mathcal{B}| - 1)$.

If (1) is the outcome for every input H , then \mathcal{A} is the required algorithm, so suppose (2) holds for some H . Then D is a set of edges of H such that every edge of $E_{S(\mathcal{B})}$ intersects $T = \bigcup_{e \in D} e \cap B$. For $u = |D \setminus V(K)|$ it is clear that

$$\left| \bigcup_{e \in D \setminus V(K)} e \cap B \right| \leq u(r - 1).$$

Next we estimate $|\bigcup_{e \in V(K)} e \cap B|$. As in the proof of Claim 6, for each component C of the constellation K , the number of vertices of B contained in $\{e_C\} \cup L_C$ is at most

$$(r - 1) + (r - 2)|L_C| = 1 + (r - 2)|V(C)|.$$

Each component of K has at least two vertices, so the number of components is at most $|V(K)|/2 = (|D| - u)/2$. Since $|D| < (2 + \varepsilon')(|\mathcal{B}| - 1)$ we find

$$\begin{aligned} \left| \bigcup_{e \in V(K)} e \cap B \right| &\leq \sum_C (1 + (r - 2)|V(C)|) \\ &\leq |V(K)|/2 + (r - 2)|V(K)| \\ &= (2r - 3)|V(K)|/2 \\ &= (2r - 3)(|D| - u)/2 \\ &< (2r - 3)(1 + \varepsilon'/2)(|\mathcal{B}| - 1) - u(r - 3/2), \end{aligned}$$

where again the sum is over all components C of K . Therefore

$$\begin{aligned} |T| &< u(r - 1) + (2r - 3)(1 + \varepsilon'/2)(|\mathcal{B}| - 1) - u(r - 3/2) \\ &= u/2 + (2r - 3 + \varepsilon'(r - 3/2))(|\mathcal{B}| - 1) \\ &< (2r - 3 + \varepsilon'(r - 1))(|\mathcal{B}| - 1) \\ &= (2r - 3 + \varepsilon)(|\mathcal{B}| - 1), \end{aligned}$$

where in the last line we used that $u < \varepsilon'(|\mathcal{B}| - 1)$ and $\varepsilon' = \varepsilon/(r - 1)$. But then this contradicts the assumption $\tau_B(E_S) > (2r - 3 + \varepsilon)(|S| - 1)$ for $S = S(\mathcal{B})$, thus proving Theorem 7.2.

7.2 Circular chromatic index

A proper circular p/q -edge-colouring of a graph G is a colouring of the edges of G with colours in $\{0, \dots, p - 1\}$ such that the difference modulo p of the colours assigned to two adjacent edges is not in $\{-(q - 1), -(q - 2), \dots, q - 1\}$. The smallest ratio p/q for which there is a proper circular p/q -edge-colouring of G is called the circular chromatic index of G .

Kaiser, Král and Škrekovski [45] proved the following result using Theorem 1.2.

Theorem 7.3. *Let $p \in \mathbb{N}$ with $p \geq 2$ and G be a cubic bridgeless graph with girth*

$$g = \begin{cases} 2(2p)^{2p-2} & \text{if } p \geq 2 \text{ is even,} \\ 2(2p)^{2p} & \text{if } p \geq 3 \text{ is odd.} \end{cases}$$

Then G admits a proper circular $(3p + 1)/p$ -edge-colouring.

To prove Theorem 7.3, Kaiser, Král and Škrekovski proved the existence of an IT in a certain auxiliary graph G' constructed using G , p and a fixed 1-factor F of G . The vertices of G' are partitioned into m classes V_i , one for each odd cycle C_i of $G - F$, where $|V_i| = |V(C_i)|$. The maximum degree of G' is at most $(2p)^{2(p-1)}$ if p is even, and at most $(2p)^{2p}$ if p is odd. Hence by Theorem 1.1 there exists an IT in G' provided $|V_i| \geq 2(2p)^{2(p-1)}$ (respectively $|V_i| \geq 2(2p)^{2p}$) for each i . Given such an IT in G' , the authors explicitly provide the required proper circular $(3p + 1)/p$ -edge-colouring of G .

We may apply Corollary 1.5 (the algorithmic version of Theorem 1.1) to obtain an algorithmic version of Theorem 7.3, and in fact in this case no weakening of the result at all is necessary. This is because the sizes of the vertex classes V_i are exactly the lengths of the odd cycles in $G - F$, and hence girth $g \geq 2(2p)^{2(p-1)}$ (respectively $g \geq 2(2p)^{2p}$) is enough to ensure the extra one in the lower bounds on the $|V_i|$ required by Corollary 1.5. Note also that the number m of odd cycles in $G - F$ is clearly less than $|V(G)|$.

Corollary 7.4. *Let $p \in \mathbb{N}$ with $p \geq 2$ be given. Then there exists an algorithm that takes as input any cubic bridgeless graph G with girth*

$$g = \begin{cases} 2(2p)^{2p-2} & \text{if } p \geq 2 \text{ is even,} \\ 2(2p)^{2p} & \text{if } p \geq 3 \text{ is odd,} \end{cases}$$

and finds, in time polynomial in $|V(G)|$, a proper circular $(3p + 1)/p$ -edge-colouring of G .

7.3 Strong colouring

Let k and n be positive integers. Let G be a graph with n vertices and let (V_1, \dots, V_m) be a vertex partition of $V(G)$ such that $|V_i| \leq k$ for all i . A graph G is *strongly k -colourable with respect to (V_1, \dots, V_m)* if there is a colouring of G with k colours so that for each vertex class, each colour is assigned to at most one vertex of the class. If G is strongly k -colourable with respect to every vertex partition of $V(G)$ into classes of size at most k , we say G is *strongly k -colourable*. The *strong chromatic number* of a graph G , denoted $s\chi(G)$, is the minimum k such that G is strongly k -colourable. This notion was introduced independently by Alon [7, 9] and Fellows [27].

The best known general bound for the strong chromatic number of graphs G in terms of their maximum degree $\Delta(G)$ is $s\chi(G) \leq 3\Delta(G) - 1$, proved in [39]. (See also [40] for an asymptotically better bound.) Aharoni, Berger and Ziv [4] gave a nice simplification of the proof in [39], that gives the bound $s\chi(G) \leq 3\Delta(G)$. Their argument uses a slight strengthening of Theorem 1.1, which states that if G is a graph with maximum degree Δ , and if $|V_i| \geq 2\Delta$ for each i , then for each vertex v there exists an IT of G containing v . This follows immediately from Theorem 1.2 applied to G with the partition $(\{v\}, V_2, \dots, V_m)$, assuming without loss of generality that $v \in V_1$. (See the note after the statement of Theorem 1.2 in the Introduction.)

To make the argument of [4] algorithmic, we just apply Theorem 1.4 instead of Theorem 1.2 in the previous paragraph. This gives the following slightly strengthened version of Corollary 1.5.

Corollary 7.5. *Let $\Delta \in \mathbb{N}$ be given. Then there exists an algorithm \mathcal{A} that takes as input any graph G with maximum degree Δ and vertex partition (V_1, \dots, V_m) such that $|V_i| \geq 2\Delta + 1$ for each i , and any $v \in V(G)$, and finds, in time polynomial in $|V(G)|$, an independent transversal in G that contains v .*

The proof in [4] begins with a partial strong 3Δ -colouring c of a graph G with respect to a vertex partition (V_1, \dots, V_m) , an uncoloured vertex v and a colour α not used by c on the vertex class of v . A new graph G' is obtained by removing from each V_i the vertices whose colour appears on the neighbourhood of the vertex w_i in V_i coloured α (if it exists). This reduces the size of each class by at most Δ . Then the strengthened version of Theorem 1.1 is used to find an IT Y of G' containing v . As shown in [4], the modification of c obtained by giving each $y_i \in Y \cap V_i$ colour α , and each w_i colour $c(y_i)$, is a partial strong 3Δ -colouring that colours more vertices than c did (in particular it colours v). Hence in at most $|V(G)|$ such steps a suitable colouring is constructed. This argument therefore gives the following corollary of Corollary 7.5.

Corollary 7.6. *Let Δ be a positive integer. There exists an algorithm \mathcal{A} that takes as input any graph G with maximum degree Δ and vertex partition (V_1, \dots, V_m) where $|V_i| \leq 3\Delta + 1$ for each i , and finds, in time polynomial in $|V(G)|$, a strong $(3\Delta + 1)$ -colouring of G with respect to (V_1, \dots, V_m) .*

7.4 Hitting sets for maximum cliques

It was first shown by Rabern [56], and later (with a best possible bound) by King [47], that when the maximum degree and clique number of a graph are close enough, the graph contains an independent set meeting all maximum cliques. Finding such a set is important for various colouring problems (see e.g. [47, 57]). King’s result is as follows, where $\omega(G)$ denotes the clique number of the graph G .

Theorem 7.7. *Let G be a graph of maximum degree Δ such that $\omega(G) > \frac{2}{3}(\Delta + 1)$. Then G contains an independent set meeting every maximum clique.*

As shown in [47], if $\omega(G) > \frac{2}{3}(\Delta + 1)$ then the set of all maximum cliques in G can be partitioned into classes such that for each class \mathcal{C}_i we have $|\cap_{C \in \mathcal{C}_i} V(C)| \geq (\Delta + 1)/3$. Thus for each class \mathcal{C}_i there exists a ‘core’ V_i of at least $(\Delta + 1)/3$ vertices that is contained in every clique in \mathcal{C}_i . Moreover, the V_i are all disjoint. Thus an IT of the subgraph of G induced by the union of all the cores V_i (where the V_i are the vertex classes) provides a suitable independent hitting set. By following the argument in [47] (and being slightly more careful with divisibility), it can be shown that it suffices to establish the following modification of Theorem 1.2 with the value $k = \lceil (\Delta + 1)/3 \rceil$.

Lemma 7.8. *Let k be a positive integer and let G be a graph with vertex partition (V_1, \dots, V_m) . If, for every i and every $v \in V_i$, the vertex v has at most $\min\{k - 1, |V_i| - k\}$ neighbours outside V_i , then G has an IT.*

In [47] the slightly stronger statement with $\min\{k - 1, |V_i| - k\}$ replaced by $\min\{k, |V_i| - k\}$ is established and used, but the proof is not algorithmic. The advantage of Lemma 7.8 is that we can use Theorem 1.4 to give an algorithmic proof, as follows.

Proof. As before, we may assume each V_i is independent. Since each vertex has at most $k - 1$ neighbours, we have that G is k -claw-free. Let $\varepsilon = 1/(k - 1)$. We apply the algorithm FindITorBD to G and (V_1, \dots, V_m) . Since k is fixed, the running time is polynomial in $|V(G)|$, where the degree of the polynomial depends only on k . We obtain either

- (1) an IT of G , or
- (2) a set \mathcal{B} of vertex classes and a set D of vertices of G such that D dominates $G_{\mathcal{B}}$ in G and $|D| < (2 + \varepsilon)(|\mathcal{B}| - 1)$. Moreover, D contains $V(K)$ for a constellation K of some $\mathcal{B}_0 \supseteq \mathcal{B}$, where $|D \setminus V(K)| < \varepsilon(|\mathcal{B}| - 1)$.

If (1) is the output for every input G then we have the required algorithm, so suppose (2) holds for some G . Recall from Definition 1.3 that the set of leaves in the constellation K forms an IT of $|\mathcal{B}_0| - 1$ vertex classes of \mathcal{B}_0 , and hence in particular D contains an IT Y of a set \mathcal{B}' of $|\mathcal{B}| - 1$ vertex classes of \mathcal{B} .

Since D dominates $G_{\mathcal{B}}$, we know that

$$\sum_{v \in D} d(v) \geq \sum_{V_i \in \mathcal{B}} |V_i| \geq \sum_{V_i \in \mathcal{B}'} |V_i|.$$

Since Y is an IT of \mathcal{B}' , we have that $|Y| \geq |\mathcal{B}'| - 1$ and so $|D \setminus Y| < (1 + \varepsilon)(|\mathcal{B}'| - 1)$. For any vertex $u \in D \setminus Y$, we know $d(u) \leq k - 1$ and for any vertex $v \in Y$, we have $d(v) \leq |V_i| - k$ where

$v \in V_i$. Hence

$$\begin{aligned} \sum_{v \in D} d(v) &\leq \sum_{v \in Y} d(v) + (k - 1)|D \setminus Y| \\ &< \sum_{V_i \in \mathcal{B}'} (|V_i| - k) + (k - 1)(1 + \varepsilon)(|\mathcal{B}| - 1) \\ &\leq \sum_{V_i \in \mathcal{B}'} |V_i| - (|\mathcal{B}| - 1) + \varepsilon(k - 1)(|\mathcal{B}| - 1) \\ &= \sum_{V_i \in \mathcal{B}'} |V_i|. \end{aligned}$$

This is a contradiction and so outcome (2) never occurs. This completes the proof. □

When Δ is fixed, the sets V_i can be found algorithmically in time polynomial in $|V(G)|$. Thus we have the following.

Theorem 7.9.

- (1) Let k be a positive integer. There exists an algorithm \mathcal{A} that finds an IT in any graph G with vertex partition (V_1, \dots, V_m) with the property that, for each i and each $v \in V_i$, the vertex v has at most $\min\{k - 1, |V_i| - k\}$ neighbours outside V_i . The running time of \mathcal{A} is polynomial in $|V(G)|$.
- (2) Let Δ be a positive integer. There exists an algorithm \mathcal{A}' that finds, in time polynomial in $|V(G)|$, an independent set meeting every maximum clique in any graph G with maximum degree Δ and $\omega(G) > 2(\Delta + 1)/3$.

8. Concluding remarks

As noted in the Introduction, the running time of the algorithm in Theorem 1.4 is polynomial in $|V(G)|$, where the degree of the polynomial depends on the input parameters ε and r . (Similarly in Corollary 1.5 it depends on Δ .) While the dependence on ε seems unavoidable, we are not certain of the nature (or even the necessity) of the dependence on r . If this could be avoided, or even if the condition of being r -claw-free could be substantially weakened, then algorithmic versions of many more applications of Theorems 1.1 and 1.2 would follow. The r -claw-free condition was required by our choice of signature vector, and used in the bounds in Lemma 5.3. We consider it an interesting open question as to whether this condition is essential. Recently, a randomized algorithm (that uses FindITorBD as a subroutine) was developed in [30] that overcomes the exponential dependence on Δ in Corollary 1.5.

See [29] for full details on the applications outlined in Section 7 as well as others (e.g. special graph partitions as in [12], and more specific results on circular chromatic index as in [45]), and more discussion of other applications of Theorems 1.1 and 1.2.

We close by pointing out that Theorems 1.1 and 1.2 have both a *combinatorial* proof [36] and a *topological* proof [5, 38, 51], using the notion of *topological connectedness* (see e.g. [3]). The algorithms presented in this paper, and also that of [14, 15, 16], are based on the combinatorial proofs of these results. There are other criteria guaranteeing the existence of independent transversals for which only a topological proof is known, for example [3, 6], which also have many applications. Thus algorithmic versions of these results would also be very interesting and useful, but currently seem out of reach. One simply stated example concerns 3-partite 3-uniform hypergraphs, in which the vertex set has a partition into 3 parts A, B and C , and each edge contains exactly one vertex from each part. (Thus a 3-partite 3-uniform hypergraph is also bipartite in the

sense of Section 7.1.) The following is a slight reformulation of the main result of [2], which is a direct application of [6].

Theorem 8.1. *Let H be a 3-partite 3-uniform hypergraph and let k be a non-negative integer. Then H contains either a set of k disjoint edges, or a set W of vertices such that $|W| \leq 2(k-1)$ and $W \cap e \neq \emptyset$ for every edge e of H .*

If the requirement $|W| \leq 2(k-1)$ is relaxed to $|W| \leq 3(k-1)$ then a simple greedy matching procedure gives an algorithmic proof. The best known algorithmic result for this problem is given by the proof in [37], which is also combinatorial, and provides a suitable algorithm for $|W| \leq 5(k-1)/2$.

Acknowledgements

The authors are indebted to Chidambaram Annamalai for very helpful discussions and correspondence. They would also like to thank Nikhil Bansal and David Harris for helpful comments.

References

- [1] Achlioptas, D. and Iliopoulos, F. (2016) Random walks that find perfect objects and the Lovász Local Lemma. *J. Assoc. Comput. Mach.* **63** 22.
- [2] Aharoni, R. (2001) Ryser's conjecture for tripartite 3-graphs. *Combinatorica* **21** 1–4.
- [3] Aharoni, R. and Berger, E. (2006) The intersection of a matroid and a simplicial complex. *Trans. Amer. Math. Soc.* **358** 4895–4917.
- [4] Aharoni, R., Berger, E. and Ziv, R. (2007) Independent systems of representatives in weighted graphs. *Combinatorica* **27** 253–267.
- [5] Aharoni, R., Chudnovsky, M. and Kotlov, A. (2002) Triangulated spheres and colored cliques. *Discrete Comput. Geom.* **28** 223–229.
- [6] Aharoni, R. and Haxell, P. (2000) Hall's theorem for hypergraphs. *J. Graph Theory* **35** 83–88.
- [7] Alon, N. (1988) The linear arboricity of graphs. *Israel J. Math.* **62** 311–325.
- [8] Alon, N. (1991) A parallel algorithmic version of the local lemma. In *Thirty-Second Annual Symposium on Foundations of Computer Science*, pp. 586–593, IEEE.
- [9] Alon, N. (1992) The strong chromatic number of a graph. *Random Struct. Algorithms* **3** 1–7.
- [10] Alon, N. (2003) Problems and results in extremal combinatorics I. *Discrete Math.* **273** 31–53.
- [11] Alon, N. and Asodi, V. (2007) Edge colouring with delays. *Combin. Probab. Comput.* **16** 173–191.
- [12] Alon, N., Ding, G., Oporowski, B. and Vertigan, D. (2003) Partitioning into graphs with only small components. *J. Combin. Theory Ser. B* **87** 231–243.
- [13] Alon, N. and Spencer, J. (2008) *The Probabilistic Method*, third edition, Wiley.
- [14] Annamalai, C. (2016) Finding perfect matchings in bipartite hypergraphs. In *Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1814–1823.
- [15] Annamalai, C. (2017) Algorithmic advances in allocation and scheduling. PhD dissertation, ETH Zurich.
- [16] Annamalai, C. (2018) Finding perfect matchings in bipartite hypergraphs. *Combinatorica* **38** 1285–1307.
- [17] Asadpour, A., Feige, U. and Saberi, A. (2008) Santa Claus meets hypergraph matchings. In *APPROX 2008 and RANDOM 2008* (A. Goel *et al.*, eds), Vol. 5171 of Lecture Notes in Computer Science, pp. 10–20, Springer.
- [18] Asadpour, A., Feige, U. and Saberi, A. (2012) Santa Claus meets hypergraph matchings. *ACM Trans. Algorithms* **8** 24.
- [19] Beck, J. (1991) An algorithmic approach to the Lovász Local Lemma. *Random Struct. Algorithms* **2** 343–365.
- [20] Bissacot, R., Fernández, R., Procacci, A. and Scoppola, B. (2011) An improvement of the Lovász local lemma via cluster expansion. *Combin. Probab. Comput.* **20** 709–719.
- [21] Bollobás, B., Erdős, P. and Szemerédi, E. (1975) On complete subgraphs of r -chromatic graphs. *Discrete Math.* **13** 97–107.
- [22] Britnell, J., Evseev, A., Guralnick, R., Holmes, P. and Maróti, A. (2008) Sets of elements that pairwise generate a linear group. *J. Combin. Theory Ser. A* **115** 442–465.
- [23] Chandrasekaran, K., Goyal, N. and Haeupler, B. (2013) Deterministic algorithms for the Lovász local lemma. *SIAM J. Comput.* **42** 2132–2155.
- [24] Christofides, D., Edwards, K. and King, A. (2013) A note on hitting maximum and maximal cliques with a stable set. *J. Graph Theory* **73** 354–360.

- [25] Czumaj, A. and Scheideler, C. (2000) Coloring non-uniform hypergraphs: a new algorithmic approach to the general Lovász local lemma. In *Eleventh Annual ACM–SIAM Symposium on Discrete Algorithms*, pp. 30–39.
- [26] Erdős, P. and Lovász, L. (1975) Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and Finite Sets (Keszthely, Hungary 1973)*, Vol. 10 of *Colloquia Mathematica Societatis János Bolyai*, pp. 609–627.
- [27] Fellows, M. (1990) Transversals of vertex partitions in graphs. *SIAM J. Discrete Math.* **3** 206–215.
- [28] Fischer, M. and Ghaffari, M. (2017) Subalgorithmic distributed algorithms for Lovász local lemma and the complexity hierarchy. In *Thirty-First International Symposium on Distributed Computing*, Vol. 91 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 18:1–18:16, Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [29] Graf, A. (2019) Finding independent transversals efficiently. PhD dissertation, University of Waterloo.
- [30] Graf, A., Harris, D. and Haxell, P. Algorithms for weighted independent transversals and strong colouring, submitted
- [31] Harris, D. (2016) Lopsidedependency in the Moser–Tardos framework: beyond the lopsided local lemma. *ACM Trans. Algorithms* **13** 17.
- [32] Harris, D. (2018) Derandomizing the Lovász local lemma via log-space statistical tests. [arXiv:1807.06672](https://arxiv.org/abs/1807.06672)
- [33] Harris, D. and Srinivasan, A. (2013) The Moser–Tardos framework with partial resampling. In *IEEE Fifty-Fourth Annual Symposium on Foundations of Computer Science*, pp. 469–478.
- [34] Harris, D. and Srinivasan, A. (2014) A constructive algorithm for the Lovász local lemma on permutations. In *Twenty-Fifth Annual ACM–SIAM Symposium on Discrete Algorithms*, pp. 907–925.
- [35] Harris, D. and Srinivasan, A. (2017) A constructive Lovász local lemma for permutations. *Theory Comput.* **13** 1–41.
- [36] Haxell, P. (1995) A condition for matchability in hypergraphs. *Graphs Combin.* **11** 245–248.
- [37] Haxell, P. (1995) A note on a conjecture of Ryser. *Periodica Math. Hungar.* **30** 73–79.
- [38] Haxell, P. (2001) A note on vertex list colouring. *Combin. Probab. Comput.* **10** 345–347.
- [39] Haxell, P. (2004) On the strong chromatic number. *Combin. Probab. Comput.* **13** 857–865.
- [40] Haxell, P. (2008) An improved bound for the strong chromatic number. *J. Graph Theory* **58** 148–158.
- [41] Haxell, P. (2011) On forming committees. *Amer. Math. Monthly* **118** 777–788.
- [42] Haxell, P. and Szabó, T. (2006) Odd independent transversals are odd. *Combin. Probab. Comput.* **15** 193–211, .
- [43] Haxell, P., Szabó, T. and Tardos, G. (2003) Bounded size components: partitions and transversals. *J. Combin. Theory Ser. B* **88** 281–297.
- [44] Jin, G. (1992) Complete subgraphs of r -partite graphs. *Combin. Probab. Comput.* **1** 241–250.
- [45] Kaiser, T., Král, D. and Škrekovski, R. (2004) A revival of the girth conjecture. *J. Combin. Theory Ser. B* **92** 41–53.
- [46] Kaiser, T., Král, D., Škrekovski, R. and Zhu, X. (2007) The circular chromatic index of graphs of high girth. *J. Combin. Theory Ser. B* **97** 1–13.
- [47] King, A. (2011) Hitting all maximum cliques with a stable set using lopsided independent transversals *J. Graph Theory* **67** 300–305.
- [48] Kolipaka, K. and Szegedy, M. (2011) Moser and Tardos meet Lovász. In *Forty-Third Annual ACM Symposium on Theory of Computing*, pp. 235–244.
- [49] Kolipaka, K., Szegedy, M. and Xu, Y. (2012) A sharper local lemma with improved applications. In *Approximation, Randomization, and Combinatorial Optimization*, Vol. 7408 of *Lecture Notes in Computer Science*, pp. 603–614, Springer.
- [50] Krivelevich, M. (1997) Almost perfect matchings in random uniform hypergraphs. *Discrete Math.* **170** 259–263.
- [51] Meshulam, R. (2001) The clique complex and hypergraph matching. *Combinatorica* **21** 89–94.
- [52] Molloy, M. and Reed, B. (1998) Further algorithmic aspects of the local lemma. In *Thirtieth Annual ACM Symposium on Theory of Computing*, pp. 524–529.
- [53] Moser, R. (2009) A constructive proof of the Lovász Local Lemma. In *Forty-First Annual ACM Symposium on Theory of Computing*, pp. 343–350.
- [54] Moser, R. and Tardos, G. (2010) A constructive proof of the general Lovász Local Lemma. *J. Assoc. Comput. Mach.* **57** 11.
- [55] Pegden, W. (2014) An extension of the Moser–Tardos algorithmic local lemma, *SIAM J. Discrete Math.* **28** 911–917.
- [56] Rabern, L. (2011) On hitting all maximum cliques with an independent set. *J. Graph Theory* **66** 32–37.
- [57] Reed, B. (2001) *Perfect Graphs*, Wiley.
- [58] Srinivasan, A. (2008) Improved algorithmic versions of the Lovász local lemma. In *Nineteenth Annual ACM–SIAM Symposium on Discrete Algorithms*, pp. 611–620.
- [59] Szabó, T. and Tardos, G. (2006) Extremal problems for transversals in graphs with bounded degree. *Combinatorica* **26** 333–351.
- [60] Yuster, R. (1997) Independent transversals in r -partite graphs. *Discrete Math.* **176** 255–261.