

Algorithms for weighted independent transversals and strong colouring

Alessandra Graf¹, David G. Harris², and Penny Haxell*¹

¹Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON, Canada

²Department of Computer Science, University of Maryland, College Park, MD, United States

August 24, 2021

Abstract

An *independent transversal* (IT) in a graph with a given vertex partition is an independent set consisting of one vertex in each partition class. Several sufficient conditions are known for the existence of an IT in a given graph and vertex partition, which have been used over the years to solve many combinatorial problems. Some of these IT existence theorems have algorithmic proofs, but there remains a gap between the best existential bounds and the bounds obtainable by efficient algorithms.

Recently, Graf and Haxell (2018) described a new (deterministic) algorithm that asymptotically closes this gap, but there are limitations on its applicability. In this paper we develop a randomized algorithm that is much more widely applicable, and demonstrate its use by giving efficient algorithms for two problems concerning the strong chromatic number of graphs.

This is an extended version of an article which appeared in the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2021.

1 Introduction

Let $G = (V(G), E(G))$ be a graph with a partition \mathcal{V} of its vertices; the elements of \mathcal{V} are non-empty subsets of $V(G)$, which we refer to as *blocks*. For a vertex $v \in V(G)$, we let $\mathcal{V}(v)$ denote the unique block $U \in \mathcal{V}$ with $v \in U$. The *minimum blocksize* $b^{\min}(\mathcal{V})$ (or just b^{\min} if \mathcal{V} is understood) is the minimum size of any block in \mathcal{V} . We say that \mathcal{V} is *b-regular* if every block U has size exactly b .

We let $n = |V(G)|$ denote the number of vertices in G . The *neighbourhood* $N(v)$ of a vertex v is the set of vertices $u \in V(G)$ with $uv \in E(G)$. The *maximum degree* $\Delta(G)$ is the maximum number of neighbours of any vertex; again, if G is understood, we write simply Δ .

An independent set I of G is called an *independent transversal* (IT) of G with respect to \mathcal{V} if $|I \cap U| = 1$ for all $U \in \mathcal{V}$; likewise, I is a *partial independent transversal* (PIT) of G with respect to \mathcal{V} if $|I \cap U| \leq 1$ for all $U \in \mathcal{V}$. Many combinatorial problems can be formulated in terms of ITs in graphs with respect to given vertex partitions (see e.g. [25]). Various results give sufficient conditions for the existence of an IT (e.g. [2, 11, 21, 22, 4, 8]). In particular, [21, 22] showed the following:

*Partially supported by NSERC.

Theorem 1 ([21, 22]). *If G has a vertex partition with $b^{\min} \geq 2\Delta(G)$, then an IT of G exists.*

This bound is optimal, since [31] showed that blocks of size $2\Delta - 1$ are not sufficient to guarantee the existence of an IT.

There is another important extension involving weighted ITs in the setting of vertex-weighted graphs. Theorem 1 (which merely shows the existence of an IT without regard to weight) is not sufficient for these applications. For a weight function $w : V(G) \rightarrow \mathbb{R}$ and a subset $U \subseteq V(G)$, we write $w(U) = \sum_{u \in U} w(u)$ and $w^{\max}(U) = \max_{u \in U} w(u)$. We also write $w(G) = w(V(G)) = \sum_{v \in V(G)} w(v)$. We say w is *non-negative* if $w(v) \geq 0$ for all v .

Aharoni, Berger, & Ziv [1] showed the following (in a different but equivalent formulation):

Theorem 2 ([1]). *If G has a b -regular vertex partition with $b \geq 2\Delta(G)$, then for any weight function $w : V(G) \rightarrow \mathbb{R}$ there exists an IT M of G with $w(M) \geq w(G)/b$.*

The proofs of Theorem 1 and Theorem 2 are not algorithmic. There are some algorithms to efficiently return an IT given a graph G and vertex partition \mathcal{V} , some of which can handle weighted graphs (see [6, 18, 19]). These mostly rely on algorithmic versions of the Lovász Local Lemma (LLL); they typically give an IT under more stringent conditions of the form $b^{\min} \geq c\Delta$, where c is a constant strictly larger than 2. The algorithm of [18] has the condition $b^{\min} \geq 4\Delta - 1$, which is the strongest known criterion of this form.

Recently, Graf & Haxell [13] developed a new algorithm, called FindITorBD [13], to find either an IT in G or a set of blocks with a small dominating set D which has some additional properties. The algorithm uses ideas from the original proof of Theorem 1 and modifications of several key notions (including “lazy updates”) from Annamalai [6].

To describe FindITorBD, we require a few definitions. A vertex set D *dominates* another vertex set W in G if for all $w \in W$, there exists $uw \in E(G)$ for some $u \in D$. (This is also known as *strong domination* or *total domination*, but it is the only notion of domination we need so we use the simpler term.) For a subset $\mathcal{B} \subseteq \mathcal{V}$ of the vertex partition, we write $V(\mathcal{B}) = \bigcup_{U \in \mathcal{B}} U$. A *constellation* K for \mathcal{B} is a pair of disjoint vertex sets $K^{\text{centre}}, K^{\text{leaf}} \subseteq V(\mathcal{B})$ with the following properties:

- $|K^{\text{leaf}}| = |\mathcal{B}| - 1$
- K^{leaf} is a PIT of G with respect to \mathcal{V}
- Each vertex $v \in K^{\text{centre}}$ has no neighbours in K^{centre} and at least one neighbour in K^{leaf}
- Each vertex $v \in K^{\text{leaf}}$ has exactly one neighbour in K^{centre} and no neighbours in K^{leaf}

We write $V(K)$ for the vertex set $K^{\text{centre}} \cup K^{\text{leaf}}$. The induced graph on $V(K)$, denoted $G[V(K)]$, is thus a collection of stars, with centres and leaves in K^{centre} and K^{leaf} respectively. These stars are all non-degenerate in the sense that they have at least one leaf.

We state a (slightly simplified) summary of the algorithm FindITorBD as follows:

Theorem 3 ([13]). *The algorithm FindITorBD takes as input a parameter $\epsilon \in (0, 1)$ and a graph G with vertex partition \mathcal{V} and finds either:*

1. *an IT in G , or*
2. *a non-empty set $\mathcal{B} \subseteq \mathcal{V}$ and a vertex set $D \subseteq V(G)$ such that D dominates $V(\mathcal{B})$ in G and $|D| < (2 + \epsilon)(|\mathcal{B}| - 1)$. Moreover, there is a constellation K for some $\mathcal{B}_0 \supseteq \mathcal{B}$ with $V(K) \subseteq D$ and $|D \setminus V(K)| < \epsilon(|\mathcal{B}| - 1)$.*

If $\Delta(G)$ and ϵ are fixed, then the runtime is $\text{poly}(n)$.

It is easy to show that if $b^{\min}(\mathcal{V}) \geq 2\Delta(G) + 1$, then no vertex set $V(\mathcal{B})$ is dominated by a set of size less than $(2 + \frac{1}{\Delta(G)})(|\mathcal{B}| - 1)$. This leads to the following result:

Corollary 4 ([13]). *The algorithm FindITorBD takes as input a graph G and a vertex partition with $b^{\min} \geq 2\Delta + 1$ and returns an IT in G . If $\Delta \leq O(1)$, the runtime is $\text{poly}(n)$.*

This asymptotically matches the bound of Theorem 1. Thus Theorem 3 and Corollary 4 offer the possibility of new algorithmic proofs; see [13, 14] for further details and applications.

From a combinatorial point of view, the algorithm FindITorBD is nearly optimal. However, from an algorithmic point of view, it is limited by its dependence on Δ and ϵ , making it efficient only when these parameters are constant. The aim of this paper is to give a new (randomized) algorithm that overcomes this limitation, as well as extending to the setting of weighted ITs. Our main theorem is as follows.

Theorem 5. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$, a graph G with a b -regular vertex partition where $b \geq (2 + \epsilon)\Delta$, and a weight function $w : V(G) \rightarrow \mathbb{R}$, and finds an IT in G with weight at least $w(G)/b$. For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

If we disregard vertex weights, this gives the immediate corollary:

Corollary 6. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$ and a graph G with a vertex partition where $b^{\min} \geq (2 + \epsilon)\Delta$, and finds an IT of G . For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

In particular, Corollary 6 has Corollary 4 as a special case (for constant Δ , we can take $\epsilon = \frac{1}{2\Delta}$), and is also stronger than all the previous LLL-based results (since we can also take ϵ to be an arbitrary fixed constant and allow Δ to vary freely).

The overall construction has three phases. In the first phase, in Section 2, we develop an algorithm FindWeightIT which is an initial attempt to achieve Theorem 5. This uses a streamlined and algorithmic version of a construction of Aharoni, Berger, and Ziv [1], overcoming some technical challenges in the analysis stemming from the fact that constellations provided by FindITorBD are slightly “defective” and only approximately dominate parts of the graph G , as compared to the non-constructive combinatorial bounds. On its own, this algorithm has two severe limitations: while its runtime is polynomial in n , it is exponential in both the blocksize b and the number of bits of precision used to specify the weight function w .

In the second phase, discussed in Section 3, we use an algorithmic version of the LLL from Moser & Tardos [29] to sparsify the graph. Given a vertex partition with blocksize $b \geq (2 + \epsilon)\Delta$, where ϵ is an arbitrary constant, this effectively reduces the blocksize b and the degree Δ to constant values, at which point FindWeightIT can be used. Unfortunately, a number of error terms accumulate in this process, including concentration losses from the degree reduction and quantization errors from the FindWeightIT algorithm. As a consequence, this only gives an IT of weight $(1 - \lambda)w(G)/b$, where λ is an arbitrarily small constant.

In the third phase, carried out in Section 4, we overcome this limitation by “oversampling” the high-weight vertices, giving the final result of Theorem 5. For maximum generality, we analyse this in terms of a linear programming (LP) formulation related to a construction of [1].

In Section 5, we demonstrate the use of Theorem 5 by providing algorithms for finding ITs which avoid a given set of vertices L as long as $|L| < b^{\min}$. Such ITs are used in a number of constructions [30, 26], many of which do not themselves overtly involve the use of weighted ITs.

In Section 6, we consider strong colouring of graphs. For a positive integer k , we say that a graph G is *strongly k -colourable with respect to vertex partition \mathcal{V}* if there is a proper vertex colouring of G with k colours so that no two vertices in the same block receive the same colour. The *strong chromatic number* of G , denoted $s\chi(G)$, is the minimum k such that G is strongly k -colourable with respect to every vertex partition of $V(G)$ into blocks of size k .

This notion was introduced independently by Alon [2, 3] and Fellows [11] and has been widely studied [12, 28, 23, 7, 1, 27, 24, 26]. The best currently-known explicit bound for strong chromatic number in terms of maximum degree is $s\chi(G) \leq 3\Delta(G) - 1$, proved in [23]. (See also [24] for an asymptotically better bound.) It is conjectured (see e.g. [1, 31]) that the correct general bound is $s\chi(G) \leq 2\Delta(G)$. There is a natural notion of *fractional strong chromatic number* (see Section 6), for which the corresponding fractional version of this conjecture was shown in [1].

Graf & Haxell [13] used FindITorBD to develop an algorithm for strong colouring with $3\Delta + 1$ colours, but, as before, the algorithm is efficient only when Δ is constant. In Section 6 we use Theorem 5 for an efficient algorithm for strong colouring with $(3 + \epsilon)\Delta$ colours for any fixed ϵ , with no restrictions on Δ . We also give an algorithmic version of the fractional strong colouring result.

We remark that if our goal was solely to show Corollary 6, then the first and third phase of the proof of Theorem 5 could be completely omitted, and the second phase could use a simpler version of the LLL. However, Corollary 6 is not enough for our applications such as strong colouring.

We also remark that the sparsification step (Phase 2) is the only part of the overall algorithm that requires randomization. It is possible to derandomize the Moser-Tardos algorithm in this setting, giving fully deterministic algorithms for weighted ITs. This requires significant technical analysis of the Moser-Tardos algorithm beyond the scope of this paper; see [17] for further details.

2 FindWeightIT

Our starting point is a procedure FindWeightIT to find a weighted IT using FindITorBD as a subroutine. This takes as input a graph G with a b -regular vertex partition \mathcal{V} and a weight function w on G . It is defined as follows:

```

1: function FINDWEIGHTIT( $G, \mathcal{V}, w$ )
2:   Set  $W := \{v \in V(G) : w(v) = w^{\max}(\mathcal{V}(v))\}$  and  $\mathcal{W} := \{U \cap W : U \in \mathcal{V}\}$ .
3:   Apply FindITorBD to graph  $G[W]$ , vertex partition  $\mathcal{W}$  and parameter  $\epsilon = \frac{1}{8b^2}$ .
4:   if FindITorBD( $G[W], \mathcal{W}, \epsilon$ ) returns an IT  $M'$  then
5:     return  $M := M'$ 
6:   else FindITorBD( $G[W], \mathcal{W}, \epsilon$ ) returns  $\mathcal{B} \subseteq \mathcal{W}$  and  $D \subseteq W$  containing constellation  $K$ 
7:     for all  $v \in V(G)$  do  $w'(v) := w(v) - |N(v) \cap D|$ 
8:     Recursively call  $M := \text{FindWeightIT}(G, \mathcal{V}, w')$ .
9:     Set  $Y := \{v \in K^{\text{leaf}} \cap V(\mathcal{B}) : |N(v) \cap D| = 1\}$ .
10:    while there is some vertex  $v \in Y \setminus M$  with  $N(v) \cap M = \emptyset$  do
11:      Choose such a vertex  $v$  arbitrarily.
12:      Update  $M \leftarrow (M \cup \{v\}) \setminus (M \cap \mathcal{V}(v))$ .
13:    return  $M$ 

```

For a weight function $w : V(G) \rightarrow \mathbb{R}$ on G , we define $|w| = \sum_{v \in V(G)} |w(v)|$, where $|w(v)|$ is the absolute value of $w(v)$. The main result we will show for this algorithm is the following:

Theorem 7. *For an integer-valued weight function w and a b -regular vertex partition \mathcal{V} with $b > 2\Delta(G)$, the algorithm FindWeightIT(G, \mathcal{V}, w) returns an IT in G of weight at least $w(G)/b$. For fixed b , its runtime is $\text{poly}(n, |w|)$.*

Before we prove Theorem 7, we note that a simple quantization step can extend FindWeightIT to handle real-valued weight functions, with a small loss in the weight of the resulting IT.

Lemma 8. *There is an algorithm that takes as inputs a parameter $\eta > 0$, a graph G with a non-negative weight function w and a b -regular vertex partition where $b > 2\Delta(G)$, and finds an IT in G of weight at least $(1 - \eta)\frac{w(G)}{b}$. For fixed b , the runtime is $\text{poly}(n, 1/\eta)$.*

Proof (assuming Theorem 7). If $w(G) = 0$, then $w(v) = 0$ for all vertices v and so w is integral; in this case we can apply Theorem 7 directly. So suppose that $w(G) > 0$. Define $\alpha = \frac{n}{\eta w(G)}$, and define a new weight function $w' : V(G) \rightarrow \mathbb{Z}_{\geq 0}$ by $w'(v) = \lfloor \alpha w(v) \rfloor$ for each v . Apply Theorem 7 to G and w' ; since $|w'| \leq \alpha w(G) = n/\eta$, the runtime is $\text{poly}(n, 1/\eta)$ for fixed b . This generates an IT M of G with $w'(M) \geq w'(G)/b$. Here

$$w'(G) = \sum_{v \in V(G)} \lfloor \alpha w(v) \rfloor > \sum_{v \in V(G)} (\alpha w(v) - 1) = \alpha w(G) - n.$$

Therefore

$$\sum_{v \in M} w(v) \geq \sum_{v \in M} \frac{w'(v)}{\alpha} = \frac{w'(M)}{\alpha} \geq \frac{w'(G)}{b\alpha} > \frac{w(G)}{b} - \frac{n}{b\alpha} = \frac{w(G)}{b}(1 - \eta). \quad \square$$

In the remainder of the section, we will prove Theorem 7.

When we call FindWeightIT(G, \mathcal{V}, w), it generates a series of recursive calls on the same graph G and vertex partition \mathcal{V} , but different weight functions $w^{(i)}$, where $w^{(0)} = w$ and $w^{(i+1)}$ is obtained from $w^{(i)}$ according to Line 7, i.e. $w^{(i+1)} = (w^{(i)})'$. To simplify notation, we assume throughout that we have fixed a graph G and a b -regular vertex partition \mathcal{V} where $b > 2\Delta$.

Proposition 9. *In line 7 we have $w'(G) > w(G) - (b - 1)(1 + \frac{1}{16b^2})(|\mathcal{B}| - 1)$.*

Proof. In order to reach line 7, FindITorBD($G[W]; \mathcal{W}$) must return a non-empty set $\mathcal{B} \subseteq \mathcal{W}$ of blocks and vertex set D dominating $V(\mathcal{B})$ with $|D| < (2 + \epsilon)(|\mathcal{B}| - 1)$. We can compute $w'(G)$ as:

$$w'(G) = \sum_{v \in V} w'(v) = \sum_{v \in V} (w(v) - |N(v) \cap D|) = w(G) - \sum_{v \in V} |N(v) \cap D|.$$

In turn, we bound this as

$$\sum_{v \in V} |N(v) \cap D| = \sum_{x \in D} |N(x)| \leq |D|\Delta < (2 + \epsilon)(|\mathcal{B}| - 1)\Delta.$$

Since $\epsilon = \frac{1}{8b^2}$ and $b \geq 2\Delta + 1$, we thus have

$$\sum_{v \in V} |N(v) \cap D| < (2 + \frac{1}{8b^2})(|\mathcal{B}| - 1)(\frac{b-1}{2}) = (b - 1)(1 + \frac{1}{16b^2})(|\mathcal{B}| - 1). \quad \square$$

Lemma 10. *For fixed b , FindWeightIT terminates in time $\text{poly}(n, |w|)$.*

Proof. We will show algorithm termination using a potential function Φ on weights w , defined as

$$\Phi(w) = -w(G) + b \sum_{U \in \mathcal{V}} w^{\max}(U) = \sum_{v \in V} (w^{\max}(\mathcal{V}(v)) - w(v)) \geq 0$$

In each iteration where FindWeightIT reaches line 8, we have

$$\Phi(w) - \Phi(w') = b \sum_{U \in \mathcal{V}} (w^{\max}(U) - w'^{\max}(U)) - (w(G) - w'(G)).$$

As D dominates $V(\mathcal{B})$, each vertex $v \in V(\mathcal{B})$ has $w'(v) \leq w(v) - 1$. By definition of W , this implies that $w'^{\max}(U) < w^{\max}(U)$ for $U \in \mathcal{B}$. So $\sum_{U \in \mathcal{V}} (w^{\max}(U) - w'^{\max}(U)) \geq |\mathcal{B}|$. Combined with Proposition 9, this shows

$$\Phi(w) - \Phi(w') \geq b|\mathcal{B}| - (b-1)\left(1 + \frac{1}{16b^2}\right)(|\mathcal{B}| - 1) = b + \left(1 - \frac{b-1}{16b^2}\right)(|\mathcal{B}| - 1) \geq b.$$

Thus, $\Phi(w^{(i+1)}) \leq \Phi(w^{(i)}) - b$ in each iteration i . Since $\Phi(w^{(i)}) \geq 0$ always, this implies that the total number of recursive calls starting from w is at most $\Phi(w)/b \leq 2|w|$.

Next, let us check that each subproblem on weight function $w^{(i)}$ runs in $\text{poly}(n, |w|)$ time. The entries of $w^{(i)}$ are changed from $w = w^{(0)}$ by at most Δi , and so arithmetic operations take $\text{poly}(n, |w|)$ time. FindITorBD runs in $\text{poly}(n)$ time since $\Delta < b/2$ and $\epsilon = \frac{1}{8b^2}$ and b is fixed. Finally, each execution of line 12 moves a vertex of Y into M , and the vertex that gets removed from that block was not in Y because Y is a PIT and so the block contains at most one element of Y . Thus each iteration of the loop increases $|Y \cap M|$ by one, so it terminates within n iterations. \square

Lemma 11. *The set M returned by FindWeightIT is an IT of G with respect to \mathcal{V} .*

Proof. We show this by strong induction on the runtime of FindWeightIT. If FindITorBD($G[W], \mathcal{W}, \epsilon$) returns an IT, then M is defined in line 5 to be this same IT. Since $G[W]$ is an induced subgraph and \mathcal{W} is the restriction of \mathcal{V} to W , this is also an IT of G with respect to \mathcal{V} .

Otherwise, when FindITorBD($G[W], \mathcal{W}, \epsilon$) returns a set \mathcal{B} of blocks and a set D of vertices, FindWeightIT is recursively applied to obtain a set M (line 8). The runtime on this recursive subproblem is clearly less than the runtime of the overall algorithm itself. So by the induction hypothesis, M is an IT with respect to \mathcal{V} .

The set M remains a transversal throughout the loop at line 10 because line 12 adds a vertex $v \in Y \setminus M$ to M and removes the vertex of M in the same block as v . Also, M remains an independent set since $N(v) \cap M = \emptyset$. Thus M at the end is an IT of G with respect to \mathcal{V} . \square

Thus FindWeightIT terminates quickly and returns an IT. It remains to show that the resulting IT M has high weight. We first show a few preliminary results.

Proposition 12. *The value of $w'(M)$ does not decrease during any iteration of the loop at line 10.*

Proof. Let v be the vertex chosen in line 11, let $U = \mathcal{V}(v)$ be the block containing v and let $a = w^{\max}(U)$. By the definition of Y (line 9), we know that $v \in W$ and that v has exactly one neighbour in D . Thus $w(v) = a$ and $w'(v) = w(v) - 1 = a - 1$.

Line 12 updates the transversal M by adding v and removing the vertex x currently in $M \cap U$. If $x \in W$, then, since $U \cap W \in \mathcal{B}$ and D dominates $V(\mathcal{B})$, this means that x has at least one neighbour in D , which implies that $w'(x) \leq w(x) - 1 \leq a - 1$. Otherwise, if $x \notin W$, then $w(x) < a$. Since w is integer-valued, this implies that $w(x) \leq a - 1$, and so $w'(x) \leq w(x) \leq a - 1$.

In either case, $w'(x) \leq w'(v)$, and so replacing x by v in M does not decrease $w'(M)$. \square

Proposition 13. *If FindWeightIT reaches line 6, then the output M of FindWeightIT satisfies*

$$\sum_{v \in M} |N(v) \cap D| > \left(1 - \frac{1}{16b}\right) (|\mathcal{B}| - 1).$$

Proof. Because of the termination condition of the loop at line 10, each vertex $v \in Y \setminus M$ has a neighbor u in M . Since Y is a PIT, we know that $u \in M \setminus Y$. So there are at least $|Y \setminus M|$ edges from $M \setminus Y$ to $Y \setminus M$. Since $Y \subseteq D$, this in turn shows that $\sum_{v \in M \setminus Y} |N(v) \cap D| \geq |Y \setminus M|$. Also, since D dominates $V(\mathcal{B})$ and $Y \subseteq V(\mathcal{B})$, any vertex $v \in M \cap Y$ has $|N(v) \cap D| \geq 1$, and hence $\sum_{v \in M \cap Y} |N(v) \cap D| \geq |M \cap Y|$. Putting these bounds together, we have

$$\sum_{v \in M} |N(v) \cap D| = \sum_{v \in M \setminus Y} |N(v) \cap D| + \sum_{v \in M \cap Y} |N(v) \cap D| \geq |Y \setminus M| + |Y \cap M| = |Y|.$$

To complete the proof, we will show that $|Y| > (1 - \frac{1}{16b})(|\mathcal{B}| - 1)$. To see this, note that

$$|Y| \geq |K^{\text{leaf}}| - |K^{\text{leaf}} \setminus V(\mathcal{B})| - |\{y \in K^{\text{leaf}} : |N(y) \cap D| \neq 1\}|.$$

Here K is a constellation for some $\mathcal{B}_0 \supseteq \mathcal{B}$ and thus $|K^{\text{leaf}}| = |\mathcal{B}_0| - 1$. Since K^{leaf} is a PIT, we have $|K^{\text{leaf}} \setminus V(\mathcal{B})| \leq |\mathcal{B}_0| - |\mathcal{B}|$. Each vertex in K^{leaf} has zero neighbours in K^{leaf} and exactly one neighbour in K^{centre} . Thus, if $y \in K^{\text{leaf}}$ has more than one neighbour in D , then it has a neighbour in $J = D \setminus V(K)$. (Recall that $V(K) \subseteq D$.) Theorem 3 ensures that $|J| < \epsilon(|\mathcal{B}| - 1)$ so there are fewer than $\epsilon\Delta(|\mathcal{B}| - 1)$ edges from J to K^{leaf} and hence fewer than $\epsilon\Delta(|\mathcal{B}| - 1)$ vertices $y \in K^{\text{leaf}}$ with $|N(y) \cap D| \neq 1$. As $\epsilon = \frac{1}{8b^2}$ and $\Delta < b/2$, we have

$$|Y| > (|\mathcal{B}_0| - 1) - (|\mathcal{B}_0| - |\mathcal{B}|) - (\frac{1}{8b^2})(b/2)(|\mathcal{B}| - 1) = (1 - \frac{1}{16b})(|\mathcal{B}| - 1). \quad \square$$

We are now ready to prove that the IT returned by FindWeightIT has the desired weight.

Lemma 14. *For $M = \text{FindWeightIT}(G, \mathcal{V}, w)$, we have $w(M) \geq w(G)/b$.*

Proof: We prove this by strong induction on the runtime of FindWeightIT. If FindITorBD returns an IT M' on the vertex set W , then

$$w(M') = \sum_{U \in \mathcal{V}} w^{\max}(U) \geq \sum_{U \in \mathcal{V}} \frac{1}{b} \sum_{v \in U} w(v) = w(G)/b,$$

and we are done.

Otherwise, suppose FindITorBD returns \mathcal{B} and D (i.e. lines 6–12 are executed). By Lemma 11, the recursive call FindWeightIT(G, \mathcal{V}, w') returns an IT M at line 8. By the induction hypothesis, it satisfies $w'(M) \geq w'(G)/b$. By Proposition 12, the value $w'(M)$ does not decrease during the loop at line 10, so the final output M also has $w'(M) \geq w'(G)/b$. By Proposition 9, we have $w'(G) > w(G) - (b-1)(1 + \frac{1}{16b^2})(|\mathcal{B}| - 1)$, and so

$$w'(M) > \frac{w(G) - (b-1)(1 + \frac{1}{16b^2})(|\mathcal{B}| - 1)}{b} = \frac{w(G)}{b} - (1 + \frac{1}{16b^2})(1 - \frac{1}{b})(|\mathcal{B}| - 1).$$

By Proposition 13, we have $w(M) - w'(M) = \sum_{v \in M} |N(v) \cap D| > (1 - \frac{1}{16b})(|\mathcal{B}| - 1)$. Overall, this gives

$$\begin{aligned} w(M) &= w'(M) + (w(M) - w'(M)) \\ &> \left(\frac{w(G)}{b} - (1 + \frac{1}{16b^2})(1 - \frac{1}{b})(|\mathcal{B}| - 1) \right) + (1 - \frac{1}{16b})(|\mathcal{B}| - 1) \\ &= \frac{w(G)}{b} + \frac{15b^2 - b + 1}{16b^3}(|\mathcal{B}| - 1) \geq \frac{w(G)}{b}. \end{aligned} \quad \square$$

Theorem 7 and Lemma 8 now follow from Lemmas 10, 11, and 14.

3 Degree Reduction

The next step in the proof is to remove the condition that b is constant. Our main tool for this is the LLL, in particular the LLL algorithm of Moser and Tardos [29]. The basic idea is to use the LLL for a “degree-splitting”: we reduce the degree, the blocksize, and the total vertex weight of G by a factor of approximately half. By doing this repeatedly, we scale down the original graph to a graph with constant blocksize. At that point we use FindWeightIT.

Let us begin by reviewing the algorithm of Moser and Tardos.

Theorem 15 ([29]). *There is a randomized algorithm which takes as input a probability space Ω in k independent variables X_1, \dots, X_k along with a collection of “bad” events B_1, \dots, B_ℓ in that space, wherein each B_i is a Boolean function of a subset of the variables $\text{var}(B_i)$.*

If $\text{epd} \leq 1$, where $p = \max_i \Pr_\Omega(B_i)$ and $d = \max_i |\{j : \text{var}(B_j) \cap \text{var}(B_i) \neq \emptyset\}|$, then the algorithm has expected runtime polynomial in k and ℓ and outputs a configuration $X = (X_1, \dots, X_k)$ such that all bad-events B_i are false on X .

One additional feature of this algorithm is critical for our application to weighted ITs: the output state X produced by the Moser-Tardos algorithm has a probability distribution with nice properties [15, 20, 16]. One result of [20], which we present in a simplified form, is the following.

Theorem 16 ([20]). *Suppose the conditions of Theorem 15 are satisfied. Let E be an event in the probability space Ω which is a Boolean function of a subset of variables $\text{var}(E)$, and let r be the number of bad-events B_i with $\text{var}(E) \cap \text{var}(B_i) \neq \emptyset$, i.e., B_i can affect E . Then the probability that E holds in the output configuration X of the Moser-Tardos algorithm is at most $e^{er} \Pr_\Omega(E)$.*

Using the Moser-Tardos algorithm, we get the following degree-splitting algorithm.

Lemma 17. *There is a randomized polynomial-time algorithm that takes as input a graph G with a non-negative weight function w and a b -regular vertex partition \mathcal{V} where $b \geq 15000$. It generates an induced subgraph G' such that*

- (i) *every block $U \in \mathcal{V}$ has exactly $b' := \lfloor b/2 \rfloor$ vertices in G' ,*
- (ii) *$\Delta(G') \leq D/2 + 10\sqrt{D \log D}$ where we define $D = \max\{b/3, \Delta(G)\}$*
- (iii) *$\frac{w(G')}{b'} \geq (1 - 1/b) \frac{w(G)}{b}$.*

Proof. We will use Theorem 15, where the probability space Ω has a variable X_U for each block $U \in \mathcal{V}$; the distribution of X_U is to select a uniformly random subset $U' \subseteq U$ of size exactly b' . We will set G' to be the induced graph on vertex set $V' = \bigcup_{U \in \mathcal{V}} X_U$. This clearly satisfies property (i).

For each vertex v , we have a bad-event B_v that v has more than $s = D/2 + 10\sqrt{D \log D}$ neighbours in V' . If all events B_v are false, then property (ii) will hold. Note that any variable X_U affects an event B_v only if $N(v) \cap U \neq \emptyset$; so, X_U can affect at most $b\Delta(G)$ events.

To calculate the parameters p and d of Theorem 15, consider some vertex v with neighbours y_1, \dots, y_k . The event B_v is affected by the variable X_{U_i} for each block $U_i = \mathcal{V}(y_i)$; each X_{U_i} in turn affects at most $b\Delta(G)$ bad-events. In total, B_v affects at most $b\Delta(G)^2 \leq 3D^3$ bad-events.

We next calculate the probability of B_v . The degree of v in V' is the sum $Y = \sum_{j=1}^k Y_j$, where Y_j is the indicator that $y_j \in V'$. The random variables Y_j are negatively correlated and Y has expectation $\mathbf{E}[Y] \leq \frac{kb'}{b} \leq \frac{D+1}{2}$. Hoeffding’s inequality applies to sums of negatively correlated random variables (see, e.g., [10]), giving:

$$\Pr(Y \geq s) \leq \Pr(Y \geq \mathbf{E}[Y] + (10\sqrt{D \log D} - 1/2)) \leq e^{-2(10\sqrt{D \log D} - 1/2)^2/k},$$

and for $D \geq 5000$ and $k \leq \Delta(G) \leq D$, this is at most D^{-100} .

Thus $p \leq D^{-100}$ and $d \leq 3D^3$, and $epd \leq 1$. The Moser-Tardos algorithm generates a configuration avoiding all bad-events B_v , and the resulting graph G' satisfies (i) and (ii). It remains to analyse $w(G')$.

By Theorem 16, for any block U and fixed b' -element set $A \subseteq U$, the probability of $X_U = A$ in the algorithm output is at most e^{epr} times its probability in the original probability space Ω , where r is the number of bad-events affected by X_U . The original sampling probability is $1/\binom{b}{b'}$ and we have already seen that $r \leq b\Delta(G) \leq 3D^2$. So

$$\Pr(X_U = A) \leq \frac{e^{eD^{-100} \cdot 3D^2}}{\binom{b}{b'}} \leq \frac{e^{D^{-97}}}{\binom{b}{b'}}.$$

Consider the random variable $L = w(G) - w(G')$. Since w is non-negative, we have:

$$\begin{aligned} \mathbf{E}[L] &= \sum_{U \in \mathcal{V}} \sum_{\substack{A \subseteq U \\ |A|=b'}} \Pr(X_U = A) w(U \setminus A) \leq \frac{e^{D^{-97}}}{\binom{b}{b'}} \sum_{U \in \mathcal{V}} \sum_{\substack{A \subseteq U \\ |A|=b'}} w(U \setminus A) \\ &= \frac{e^{D^{-97}}}{\binom{b}{b'}} \sum_{U \in \mathcal{V}} w(U) \binom{b-1}{b'} = (1 - b'/b) e^{D^{-97}} w(G). \end{aligned}$$

By Markov's inequality applied to the non-negative random variable L , therefore, the bound

$$L \leq w(G)(1 - b'/b)(1 + D^{-2}) \tag{1}$$

holds with probability at least $1 - \frac{e^{D^{-97}}}{1+D^{-2}} \geq \Omega(D^{-2})$. We can repeatedly call the algorithm until we get a configuration satisfying Eq. (1). This takes $O(D^2) = \text{poly}(n)$ repetitions on average, and each iteration has expected runtime $\text{poly}(n)$. The resulting graph G' then has

$$w(G') = w(G) - L \geq w(G)(1 - (1 - b'/b)(1 + D^{-2})) = w(G)(b'/b)(1 + D^{-2} - D^{-2}b/b')$$

and thus, since $b' \geq b/2$ and $D \geq b/3 \geq 5000$, we have

$$w(G')/b' \geq (1 - D^{-2})w(G)/b \geq (1 - 1/b)w(G)/b. \quad \square$$

Lemma 18. *There is a randomized algorithm that takes as input parameters $\epsilon, \lambda \in (0, 1)$, a graph G with a b -regular vertex partition where $b \geq (2 + \epsilon)\Delta(G)$ and a non-negative weight function w . It generates an IT with weight at least $\frac{w(G)}{b}(1 - \lambda)$. For fixed ϵ and λ , the expected runtime is $\text{poly}(n)$.*

Proof. If $b \leq \frac{10^{20}}{\epsilon^3 \lambda}$, then we can simply apply Lemma 8 directly. So, let us assume that $b > \frac{10^{20}}{\epsilon^3 \lambda}$. Our strategy will be to repeatedly apply Lemma 17 for $t = \lceil \log_2 \frac{b\epsilon^3 \lambda}{10^{20}} \rceil$ rounds. This generates a series of induced graphs $G_i = G[V_i]$ for $i = 0, \dots, t$ where $V(G) = V_0 \supseteq V_1 \supseteq \dots \supseteq V_t$, along with corresponding vertex partitions $\mathcal{V}_i = \{U \cap V_i \mid U \in \mathcal{V}\}$. At the end of this process, we finish by applying Lemma 8 to the graph G_t to get the desired independent transversal.

To analyse this process, let us recursively define parameters b_i, δ_i as:

$$\begin{aligned} b_0 &= b, & \delta_0 &= \frac{b}{2 + \epsilon}, \\ b_{i+1} &= \lceil b_i/2 \rceil, & \delta_{i+1} &= \delta_i/2 + 10\sqrt{\delta_i \log \delta_i} \quad \text{for } i = 0, \dots, t-1. \end{aligned}$$

Note the following straightforward bounds for $i \leq t$:

$$b_i = \lceil 2^{-i} b \rceil \geq 2^{-t} b \geq \frac{10^{20}}{\epsilon^3 \lambda}, \quad \text{and} \quad \delta_i \geq 2^{-i} \delta_0 \geq 2^{-t} \delta_0 \geq \frac{10^{20}}{3\epsilon^3 \lambda}. \tag{2}$$

Each partition \mathcal{V}_i during this process will be b_i -regular. The precondition of Lemma 17 at each round i , namely $b_i \geq 15000$, follows immediately from Eq. (2). To explain the role of the parameter δ_i , we show the following three bounds for $i \leq t$ by induction on i :

$$\delta_i \leq \delta_0 2^{-i} + (\delta_0 2^{-i})^{2/3} \quad (3)$$

$$b_i < 3\delta_i \quad (4)$$

$$\Delta(G_i) \leq \delta_i. \quad (5)$$

The base case $i = 0$ is clear for all of them. For the induction step for Eq. (3), let $x = \delta_0 2^{-i}$. Applying the induction hypothesis $\delta_i \leq x + x^{2/3}$ gives

$$\delta_{i+1} = \delta_i/2 + 10\sqrt{\delta_i \log \delta_i} \leq (x + x^{2/3})/2 + 10\sqrt{(x + x^{2/3}) \log(x + x^{2/3})}.$$

Since $x = \delta_0 2^{-i} \geq 10^{20}/3$, it can be easily checked this is at most $x/2 + (x/2)^{2/3}$ as desired. Next, for Eq. (4), we have

$$b_{i+1} - 3\delta_{i+1} \leq (b_i/2 + 1/2) - 3(\delta_i/2 + 10\sqrt{\delta_i \log \delta_i}) = (b_i - 3\delta_i)/2 + (1/2 - 30\sqrt{\delta_i \log \delta_i}).$$

By the induction hypothesis, the first term is negative. Since $\delta_i \geq 10^{20}/3$, the second term is also negative. Thus we maintain $b_i < 3\delta_i$ for all i .

Finally, for Eq. (5), note that when applying Lemma 17 at round i , we have $\Delta(G_{i+1}) \leq D_i/2 + 10\sqrt{D_i \log D_i}$ where $D_i = \max\{b_i/3, \Delta(G_i)\}$. By the induction hypothesis, we have $\Delta(G_i) \leq \delta_i$ and $b_i/3 < \delta_i$. Thus, $D_i \leq \delta_i$ and so $\Delta(G_{i+1}) \leq \delta_i/2 + 10\sqrt{\delta_i \log \delta_i} = \delta_{i+1}$.

Now, after applying Lemma 17 in every round, the resulting weights $w(G_i)$ satisfy:

$$\frac{w(G_{i+1})}{b_{i+1}} \geq \frac{w(G_i)}{b_i} (1 - 1/b_i)$$

and using the identity $1 - x \geq e^{-2x}$ for $x \leq 1/2$, this telescopes as:

$$\frac{w(G_t)}{b_t} \geq \frac{w(G)}{b} \prod_{i=0}^{t-1} (1 - 1/b_i) \geq \frac{w(G)}{b} \prod_{i=0}^{t-1} e^{-2/b_i} = \frac{w(G)}{b} e^{-2 \sum_{i=0}^{t-1} 1/b_i}.$$

From Eq. (2), we see that $\sum_{i=0}^{t-1} 1/b_i \leq \sum_{i=0}^{t-1} 2^i/b \leq 2^t/b \leq \frac{\epsilon^3 \lambda}{10^{20}} \leq \lambda/20$. Thus,

$$\frac{w(G_t)}{b_t} \geq \frac{w(G)}{b} e^{-\lambda/10}.$$

Finally, we check the preconditions of Lemma 8 for the graph G_t , i.e., that $b_t > 2\Delta(G_t)$ and $b_t \leq O(1)$. First, from Eq. (3) and Eq. (2) we get

$$\delta_t \leq \delta_0 2^{-t} (1 + (\delta_0 2^{-t})^{-1/3}) \leq \delta_0 2^{-t} \left(1 + \left(\frac{10^{20}}{3\epsilon^3 \lambda}\right)^{-1/3}\right) \leq 2^{-t} \delta_0 (1 + \epsilon/10^6).$$

Hence, using Eq. (5), we have

$$b_t - 2\Delta(G_t) > b_t - 2\delta_t \geq 2^{-t} b - 2 \cdot 2^{-t} \delta_0 (1 + \epsilon/10^6) = 2^{-t} b \left(1 - \frac{2(1 + \epsilon/10^6)}{2 + \epsilon}\right) > 0.$$

Also, by definition of t , we have $b_t = \lceil b2^{-t} \rceil \leq 1 + 2 \cdot 10^{20}/\epsilon^3 \lambda$; since ϵ, λ are fixed, this is fixed as well. So Lemma 8 on graph G_t and parameter $\eta = \lambda/2$ produces an IT M of weight

$$w(M) \geq \frac{w(G_t)}{b_t} (1 - \eta) \geq \frac{w(G)}{b} (1 - \lambda/2) e^{-\lambda/10} \geq \frac{w(G)}{b} (1 - \lambda). \quad \square$$

4 An LP for Weighted ITs

To finish the proof of Theorem 5, we need to remove the term $1 - \lambda$ from Lemma 18. For maximum generality, we use an LP formulation adapted from [1], which relaxes the condition $b \geq (2 + \epsilon)\Delta$ to allow each vertex v to be taken with a fractional multiplicity $\gamma_v \in [0, 1]$. Formally, given a graph G , vertex partition \mathcal{V} , weight function w , and a value $\delta \in \mathbb{R}_{\geq 0}$, we define $\mathcal{P}_{G,\delta}$ to be the following LP. (Here, δ plays the role of $\frac{1}{2+\epsilon}$.)

$$\begin{aligned} \max \quad & \sum_{v \in V(G)} w(v)\gamma_v \\ \text{subject to} \quad & \sum_{u \in N(v)} \gamma_u \leq \delta \quad \forall v \in V(G) \\ & \sum_{v \in U} \gamma_v = 1 \quad \forall U \in \mathcal{V} \\ & 0 \leq \gamma_v \leq 1 \quad \forall v \in V(G). \end{aligned}$$

If the LP $\mathcal{P}_{G,\delta}$ is feasible, we let $\tau_{G,w,\delta}$ be the largest objective function value. The next results show how to get a fractional version of Theorem 5 in terms of $\mathcal{P}_{G,\delta}$.

Proposition 19. *There is a randomized algorithm which takes as input parameters $\delta, \lambda \in (0, 1/2)$, a graph G with vertex partition \mathcal{V} , a vector γ in $\mathcal{P}_{G,\delta}$, and a non-negative weight function w on G , and returns an IT in G with weight at least $(1 - \lambda) \sum_{v \in V(G)} \gamma_v w(v)$. For fixed δ and λ , the expected runtime is poly(n).*

Proof. Let $\epsilon = 1/2 - \delta$. Form a new graph G' by creating, for each vertex $v \in V(G)$, a group of $\lceil b\gamma_v \rceil$ new vertices in $V(G')$ where $b = \lceil \frac{2n}{\epsilon\lambda} \rceil$. Also, G' has an edge u_1u_2 iff $f(u_1)f(u_2) \in E(G)$, where $f : V(G') \rightarrow V(G)$ is the function mapping each vertex $u \in V(G')$ to its corresponding vertex in $V(G)$. (So G' is a blow-up of G by independent sets.) We define a vertex partition on G' by $\mathcal{V}' = \{f^{-1}(U) : U \in \mathcal{V}\}$ and a weight function w' on G' by $w'(u) = w(f(u))$.

Let us note a few bounds on G' . The weight of G' is given by

$$w'(G') = \sum_{u \in V(G')} w'(u) = \sum_{v \in V(G)} \lceil b\gamma_v \rceil w(v) \geq b \sum_{v \in V(G)} \gamma_v w(v).$$

Now consider some vertex $u \in V(G')$ with $f(u) = v \in V(G)$. Since γ satisfies the LP, we have

$$\deg(u) = \sum_{x \in N(v)} \lceil b\gamma_x \rceil \leq \sum_{x \in N(v)} (b\gamma_x + 1) \leq n + b \sum_{x \in N(v)} \gamma_x \leq n + b\delta.$$

Each block $U' = f^{-1}(U) \in \mathcal{V}'$ has size $|U'| = \sum_{x \in U} \lceil b\gamma_x \rceil$; since $\sum_{x \in U} \gamma_x = 1$, this implies that

$$b \leq |U'| \leq n + b. \tag{6}$$

Next, in light of Eq. (6), we form a graph G'' by discarding the lowest-weight $|U'| - b$ vertices in each block $U' \in \mathcal{V}'$. The resulting blocks are b -regular, and clearly $\Delta(G'') \leq \Delta(G') \leq n + b\delta$. Since w is non-negative, discarding the lowest-weight vertices gives $w'(G'') \geq (\frac{b}{n+b})w'(G')$.

We apply Lemma 18 to this graph G'' and weight function w' , with parameters $\epsilon' = \frac{\epsilon}{1-\epsilon}$ and $\lambda' = \lambda/2$ in place of ϵ, λ . The preconditions of Lemma 18 hold since w , and hence w' , is non-negative, and for $\epsilon < 1/2$ we have:

$$\frac{b}{\Delta(G'')} \geq \frac{b}{n + b\delta} = \frac{1}{n/b + \delta} \geq \frac{1}{\frac{\epsilon}{2} + (1/2 - \epsilon)} = 2(1 + \epsilon').$$

This gives an IT M'' of G'' with $w'(M'') \geq \frac{w'(G'')}{b}(1 - \frac{\lambda}{2})$. Then $M = f(M'')$ is an IT of G with weight $w(M) = w'(M'')$. Since $n/b \leq \lambda/2$, we then have:

$$w(M) \geq \frac{w'(G'')}{b} \left(1 - \frac{\lambda}{2}\right) \geq \left(\frac{b}{n+b}\right) \left(1 - \frac{\lambda}{2}\right) w'(G) \geq (1 - \lambda) \sum_{v \in V(G)} \gamma_v w(v)$$

as desired. For fixed δ and λ , the values ϵ', λ' are fixed, and the values $b, |V(G'')|$ are polynomial in n . So the expected runtime is $\text{poly}(n)$. \square

Theorem 20. *There is a randomized algorithm which takes as input a parameter $\delta \in (0, 1/2)$, a graph G with vertex partition where $\mathcal{P}_{G,\delta} \neq \emptyset$, and a weight function w on G , and returns an IT in G with weight at least $\tau_{G,w,\delta}$. For fixed δ , the expected runtime is $\text{poly}(n)$.*

Proof. Let $\epsilon = 1/2 - \delta$. We begin by sorting the vertices in each block in descending order of weight; the vertices in block U are labeled as $v_{U,1}, v_{U,2}, \dots, v_{U,t_U}$ with $w(v_{U,1}) \geq w(v_{U,2}) \geq \dots \geq w(v_{U,t_U})$. Next, we solve the LP to obtain a solution $\gamma \in \mathcal{P}_{G,\delta}$ with $\sum_{v \in V(G)} \gamma_v w(v) = \tau_{G,w,\delta}$. This takes $\text{poly}(n)$ time since $\mathcal{P}_{G,\delta}$ has $\text{poly}(n)$ constraints.

To simplify the notation, we write $\gamma_{U,j}$ and $w_{U,j}$ as short-hand for $\gamma_{v_{U,j}}$ and $w(v_{U,j})$, respectively.

Since $\sum_{v \in U} \gamma_v = 1$, each block U has a smallest index s_U with $\sum_{k=1}^{s_U} \gamma_{U,k} \geq 1 - \epsilon$. Form a new graph G' by discarding vertices $v_{U,j}$ with $j > s_U$ in each block U , and define a weight function w' on G' by $w'(v) = w(v) - w_{U,s_U}$ for $v \in U$. Because of the sorted order of the vertices, w' is non-negative. We also define a multiplicity vector $\gamma' \in [0, 1]^{V(G')}$ for G' ; again, to simplify the notation we write $\gamma'_{U,j}$ instead of $\gamma'_{v_{U,j}}$. The vector γ' is defined by

$$\gamma'_{U,j} = \begin{cases} \frac{\gamma_{U,j}}{1-\epsilon} & \text{if } j < s_U \\ 1 - \frac{\sum_{k=1}^{s_U-1} \gamma_{U,k}}{1-\epsilon} & \text{if } j = s_U. \end{cases}$$

Note that $\gamma'_{U,s_U} \geq 0$ by definition of s_U . For $j < s_U$, we have $\gamma'_{U,j} = \frac{\gamma_{U,j}}{1-\epsilon} \leq \frac{\sum_{i=0}^j \gamma_{U,i}}{1-\epsilon} \leq \frac{1-\epsilon}{1-\epsilon} \leq 1$, and clearly $\gamma'_{U,s_U} \leq 1$. We also claim that the following bound holds for all U, j :

$$\gamma'_{U,j} \leq \frac{\gamma_{U,j}}{1-\epsilon}. \quad (7)$$

It is clear for $j < s_U$, while for $j = s_U$, we have

$$\gamma'_{U,s_U} = \frac{1 - \epsilon - \sum_{k=1}^{s_U-1} \gamma_{U,k}}{1 - \epsilon} = \frac{\gamma_{U,s_U} + 1 - \epsilon - \sum_{k=1}^{s_U} \gamma_{U,k}}{1 - \epsilon} \leq \frac{\gamma_{U,s_U}}{1 - \epsilon}.$$

We next claim that $\gamma' \in \mathcal{P}_{G',\delta'}$ where $\delta' = \frac{1/2-\epsilon}{1-\epsilon} \in (0, 1/2)$. To see this, note that the constraint $\sum_k \gamma'_{U,k} = 1$ follows from the definition of γ'_{U,s_U} and the constraint $\sum_{u \in N(v)} \gamma'_u \leq \delta'$ follows from Eq. (7) and the fact that $\gamma \in \mathcal{P}_{G,\delta}$.

So let us set $\lambda' = \epsilon$ and apply Proposition 19 to $G', w', \gamma', \delta', \lambda'$ to get an IT M of G' with

$$w'(M) \geq (1 - \lambda') \sum_{v \in V(G')} \gamma'_v w'(v) = \frac{1 - \lambda'}{1 - \epsilon} \sum_{U \in \mathcal{V}} \sum_{j=1}^{s_U-1} \gamma_{U,j} (w_{U,j} - w_{U,s_U}) = \sum_{U \in \mathcal{V}} \sum_{j=1}^{s_U-1} \gamma_{U,j} (w_{U,j} - w_{U,s_U}),$$

where we omit the summand $j = s_U$ since $w'(v_{U,s_U}) = 0$. Since M is an IT of G' , we have

$$\begin{aligned} w(M) &= w'(M) + \sum_{U \in \mathcal{V}} w_{U,s_U} \geq \sum_{U \in \mathcal{V}} \left(w_{U,s_U} + \sum_{j=1}^{s_U-1} \gamma_{U,j} (w_{U,j} - w_{U,s_U}) \right) \\ &= \sum_{U \in \mathcal{V}} \left(w_{U,s_U} + \sum_{j=1}^{t_U} \gamma_{U,j} (w_{U,j} - w_{U,s_U}) - \sum_{j=s_U}^{t_U} \gamma_{U,j} (w_{U,j} - w_{U,s_U}) \right). \end{aligned}$$

Since $\sum_{j=1}^{t_U} \gamma_{U,j} = 1$ and $\sum_U \sum_{j=1}^{t_U} \gamma_{U,j} w_{U,j} = \tau_{G,w,\delta}$, this is equal to

$$\tau_{G,w,\delta} - \sum_{U \in \mathcal{V}} \sum_{j=s_U}^{t_U} \gamma_{U,j} (w_{U,j} - w_{U,s_U}).$$

Furthermore, since $w_{U,j} \leq w_{U,s_U}$ for $j > s_U$, this in turn is at least $\tau_{G,w,\delta}$.

When δ is fixed, then so are δ', λ' , and so Proposition 19 has $\text{poly}(n)$ expected runtime. \square

As a simple corollary, this gives Theorem 5 (stated again here for convenience).

Theorem 5. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$, a graph G with a b -regular vertex partition where $b \geq (2 + \epsilon)\Delta$, and a weight function $w : V(G) \rightarrow \mathbb{R}$, and finds an IT in G with weight at least $w(G)/b$. For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

Proof. Let $\delta = \frac{1}{2+\epsilon}$ and observe that $\gamma_v = \frac{1}{b}$ is a solution to $\mathcal{P}_{G,\delta}$, since every vertex v has

$$\sum_{u \in N(v)} \gamma_u = \sum_{u \in N(v)} \frac{1}{b} \leq \frac{\Delta}{b} \leq \frac{1}{2+\epsilon} = \delta,$$

and clearly $\sum_{v \in U} \gamma_v = 1$ for each block U . We thus have $\tau_{G,w,\delta} \geq \sum_v \gamma_v w(v) = w(G)/b$. Now apply Theorem 20 with parameter δ ; note that if ϵ is fixed then so is δ . \square

4.1 Weighted PITs

An LP relaxation similar to $\mathcal{P}_{G,\delta}$ can be formulated for weighted PITs. Given a graph G , vertex partition \mathcal{V} , weight function $w : V(G) \rightarrow \mathbb{R}$, and value $\delta \in \mathbb{R}_{\geq 0}$, we define $\tau_{G,w,\delta}^*$ to be the largest objective function value to the following LP denoted $\mathcal{P}_{G,\delta}^*$.

$$\begin{aligned} \max \quad & \sum_{v \in V(G)} w(v) \gamma_v \\ \text{subject to} \quad & \sum_{u \in N(v)} \gamma_u \leq \delta \quad \forall v \in V(G) \\ & \sum_{v \in U} \gamma_v \leq 1 \quad \forall U \in \mathcal{V} \\ & 0 \leq \gamma_v \leq 1 \quad \forall v \in V(G). \end{aligned}$$

Aharoni, Berger, and Ziv [1] proved the following:¹

Theorem 21 ([1, Theorem 10]). *For any weight function w , G has a PIT of weight at least $\tau_{G,w,1/2}^*$.*

This is used by [1] to show Theorem 2. Our results for weighted ITs lead to the following analogue of Theorem 21:

Corollary 22. *There is a randomized algorithm which takes as inputs a parameter $\delta \in (0, 1/2)$, a graph G with a vertex partition and a weight function w , and finds a PIT in G with weight at least $\tau_{G,w,\delta}^*$. For fixed δ , the expected runtime is $\text{poly}(n)$.*

Proof. Let G' be the graph obtained by adding, for each block U , an isolated dummy vertex x_U with weight zero. Any solution γ to $\mathcal{P}_{G,\delta}^*$ corresponds a solution to $\mathcal{P}_{G',\delta}$, by setting $\gamma_{x_U} = 1 - \sum_{v \in U} \gamma_v$. Thus, applying Theorem 20 to graph G' gives an IT M' with weight at least $\tau_{G',w,\delta} \geq \tau_{G,w,\delta}^*$. Removing the dummy vertices from M' yields a PIT M of G with $w(M) = w(M') \geq \tau_{G,w,\delta}^*$. \square

¹The statement in [1] uses very different notation, and is formulated in terms of the dual LP.

5 Independent Transversals with Vertex Restrictions

A number of combinatorial constructions use independent transversals with additional constraints. One common restriction is that the IT must include certain vertices or be disjoint from a given set of vertices. Some LLL-based algorithms for ITs can accommodate these restrictions, sometimes with additional slack in parameters [16].

Our results on weighted independent transversals give the following crisp characterization:

Theorem 23. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$, a graph G with a vertex partition \mathcal{V} where $b^{\min} \geq (2 + \epsilon)\Delta$, and a vertex set $L \subseteq V(G)$ of size $|L| < b^{\min}$. It returns an IT disjoint from L . For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

Proof. By discarding extra vertices from each block, we may assume without loss of generality that \mathcal{V} is b -regular and $b \geq (2 + \epsilon)\Delta$ and $|L| < b$. Define a weight function on $V(G)$ by $w(v) = -1$ for $v \in L$ and $w(v) = 0$ for $v \notin L$. Now use Theorem 5 to obtain an IT M with $w(M) \geq w(G)/b = -|L|/b > -1$. Since w takes only values -1 and 0 , it must be that $w(M) = 0$ and hence $M \cap L = \emptyset$. \square

As a simple corollary of Theorem 23, we can also get ITs which *include* certain given vertices.

Corollary 24. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$, a graph G with a vertex partition \mathcal{V} where $b^{\min} \geq (2 + \epsilon)\Delta$, and a pair of vertices v_1, v_2 in the same block as each other. It returns a pair of ITs M_1 and M_2 of G such that $v_1 \in M_1, v_2 \in M_2$ and $M_1 \setminus \{v_1\} = M_2 \setminus \{v_2\}$. For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

Proof. Let $U = \mathcal{V}(v_1) = \mathcal{V}(v_2)$. Then apply Theorem 23 to graph $G' = G[V(G) \setminus U]$ with associated partition $\mathcal{V}' = \mathcal{V} \setminus \{U\}$ and with $L = N(v_1) \cup N(v_2)$. Here $|L| \leq 2\Delta(G') < b^{\min}(\mathcal{V}')$ as required. This generates an IT M' of G' ; now set $M_i = M' \cup \{v_i\}$ for $i = 1, 2$. \square

By applying Corollary 24 with $v_1 = v_2$, we get the following even simpler corollary:

Corollary 25. *There is a randomized algorithm which takes as inputs a parameter $\epsilon > 0$, a graph G with a vertex partition where $b^{\min} \geq (2 + \epsilon)\Delta$, and a vertex $v \in V(G)$, and returns an IT M containing v . For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

As some examples, a construction in [26] uses a (non-algorithmic) version of Corollary 24. We will also use Corollary 25 in our application to strong colouring next. This demonstrates the power of weighted independent transversals and Theorem 5, even in contexts without an overt weight function.

6 Strong Colouring

Aharoni, Berger, and Ziv [1] showed that the strong chromatic number is at most 3Δ using an extension of Theorem 1 giving a sufficient condition for the existence of an IT containing a specified vertex. Using Corollary 25 for this instead, we obtain the following strong colouring algorithm:

Corollary 26. *There is a randomized algorithm that takes as input a graph G with a b -regular vertex partition \mathcal{V} where $b \geq (3 + \epsilon)\Delta$, and returns a strong b -colouring of G with respect to \mathcal{V} . For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

Proof. The proof is essentially the same as that of [1], so we just give a sketch. Consider a partial strong b -colouring c of G with respect to \mathcal{V} , an uncoloured vertex v , and a colour α not used by c on the block $\mathcal{V}(v)$. Define a new graph G' by removing from each block U the vertices whose colour appears on the neighbourhood of the vertex x_U in U coloured α (if it exists). This reduces the size of each block by at most $\Delta(G)$. Then we apply Corollary 25 to find an IT Y of G' containing v . As shown in [1], if we modify c by giving each vertex $y \in Y$ colour α and the corresponding vertex $x_{\mathcal{V}(y)}$ colour $c(y)$, we obtain a partial strong b -colouring with strictly fewer uncoloured vertices than c (in particular it colours v). Hence in at most n such steps we get a strong b -colouring of G . \square

Analogous to the connection between chromatic number and fractional chromatic number, there is a fractional version of strong colouring for a graph G with a b -regular vertex partition. By LP duality, this has two equivalent definitions:

- (Primal) For all weight functions $w : V(G) \rightarrow \mathbb{R}$, there is an IT M of G with $w(M) \geq w(G)/b$.
- (Dual) There exists a function f mapping each IT M of G to a real number $f(M) \in [0, 1]$ such that $\sum_M f(M) = b$ and for all vertices $v \in V(G)$ it holds $\sum_{M \ni v} f(M) = 1$.

Observe that if the function f in the dual definition takes values in the range $f(M) \in \{0, 1\}$, then f is a strong colouring with respect to the vertex partition. The fractional version of the strong colouring conjecture mentioned in Section 1 was shown by [1]:

Theorem 27 ([1]). *Every graph G is fractionally strongly $2\Delta(G)$ -colourable.*

Theorem 27, in terms of the primal definition of fractional strong colouring, is simply a restatement of Theorem 2. Theorem 5 can be viewed as an algorithmic counterpart. There is also a generic method of [9] to convert this into an algorithmic version of the dual definition. We quote the following crisp formulation of [5]:

Theorem 28 ([5]). *Suppose that \mathfrak{S} is a collection of subsets of ground set U with associated weights g_u for each $u \in U$. Suppose that there is a polynomial-time algorithm which takes as input a weight function $w : U \rightarrow \mathbb{R}_{\geq 0}$, and returns some $S \in \mathfrak{S}$ with $\sum_{u \in S} w(u) \geq \sum_{u \in U} w(u)g_u$.*

Then there is a polynomial-time algorithm to generate a subcollection $\mathfrak{S}_0 \subseteq \mathfrak{S}$ with $|\mathfrak{S}_0| \leq |U|$, with associated weights $\lambda : \mathfrak{S}_0 \rightarrow [0, 1]$, such that $\sum_{S \in \mathfrak{S}_0} \lambda(S) = 1$ and for every $u \in U$, it holds that $\sum_{S \in \mathfrak{S}_0: u \in S} \lambda(S) \geq g_u$.

As an immediate corollary of Theorem 28 and Theorem 5, we obtain the following:

Theorem 29. *There is a randomized algorithm which takes as input a parameter $\epsilon > 0$ and a graph G with b -regular vertex partition \mathcal{V} where $b \geq (2 + \epsilon)\Delta$, and finds ITs I_1, \dots, I_n with associated weights $f_1, \dots, f_n \geq 0$, such that $f_1 + \dots + f_n = b$ and for every vertex $v \in V$ it holds that $\sum_{i: v \in I_i} f_i = 1$. For fixed ϵ , the expected runtime is $\text{poly}(n)$.*

Proof. Apply Theorem 28 where \mathfrak{S} is the collection of ITs of G , and where the ground set U is $V(G)$, and where $g_v = 1/b$ for every vertex v . By Theorem 5, we have a polynomial-time procedure to find an IT S for any given weight function w with $\sum_{v \in S} w(v) \geq \sum_{v \in V} w(v)/b$. By Theorem 28, we get a collection $\mathfrak{S}_0 = \{I_1, \dots, I_n\}$ of ITs and weights λ with $\sum_{I \in \mathfrak{S}_0: v \in I} \lambda(I) \geq 1/b$ for all v and $\sum_{I \in \mathfrak{S}_0} \lambda(I) = 1$. We set $f_i = b\lambda(I_i)$ for each i . Since each $I \in \mathfrak{S}_0$ has exactly one vertex in each block, we must have $\sum_{i: v \in I_i} f_i = 1$ for all v . \square

7 Acknowledgments

Thanks to journal and conference reviewers for many helpful corrections and suggestions.

References

- [1] R. Aharoni, E. Berger, and R. Ziv. Independent systems of representatives in weighted graphs, *Combinatorica* 27(3): 253-267, 2007.
- [2] N. Alon. The linear arboricity of graphs, *Israel Journal of Mathematics* 62(3): 311-325, 1988.
- [3] N. Alon. The strong chromatic number of a graph, *Random Structures and Algorithms* 3: 1-7, 1992.
- [4] N. Alon and J. Spencer. *The Probabilistic Method*, 3rd ed., Wiley, New York, 2008.
- [5] G. Anegg, H. Angelidakis, and R. Zenklusen. Simpler and stronger approaches for non-uniform hypergraph matching and the Füredi, Kahn, and Seymour Conjecture. *Proc. 4th SIAM Symposium on Simplicity in Algorithm (SOSA)*:196-203, 2021.
- [6] C. Annamalai. Finding perfect matchings in bipartite hypergraphs, *Combinatorica*. 38(6): 1285-1307, 2018.
- [7] M. Axenovich and R. Martin. On the strong chromatic number of graphs, *SIAM Journal on Discrete Mathematics* 20(3): 741-747, 2006.
- [8] R. Bissacot, R. Fernández, A. Procacci, and B. Scoppola. An improvement of the Lovász local lemma via cluster expansion, *Combinatorics, Probability and Computing* 20(5): 709-719, 2011.
- [9] R. Carr, S. Vempala. Randomized metarounding, *Random Structures & Algorithms* 20(3): 343-352, 2002.
- [10] A. Farcomeni. Some finite sample properties of negatively dependent random variables, *Theory of Probability and Mathematical Statistics* 77: 155-163, 2008.
- [11] M. Fellows. Transversals of vertex partitions in graphs, *SIAM Journal of Discrete Mathematics* 3(2): 206-215, 1990.
- [12] H. Fleischner and M. Stiebitz. A solution to a colouring problem of P. Erdős, *Discrete Mathematics* 101(1-3): 39-48, 1992.
- [13] A. Graf and P. Haxell. Finding independent transversals efficiently, *Combinatorics, Probability and Computing* 29(5): 780-806, 2020.
- [14] A. Graf. Finding independent transversals efficiently. PhD thesis, University of Waterloo, 2019.
- [15] B. Haeupler, B. Saha, and A. Srinivasan. New constructive aspects of the Lovász local lemma, *Journal of the ACM* 58(6): Article #28, 2011.
- [16] D. Harris. New bounds for the Moser-Tardos distribution, *Random Structures & Algorithms* 57(1): 97-131, 2020.
- [17] D. Harris. Deterministic algorithms for the Lovász Local Lemma: simpler, more general, and more parallel, arXiv: 1909.08065 (2019)
- [18] D. Harris. Lopsidedependency in the Moser-Tardos framework: beyond the Lopsided Lovász Local Lemma, *ACM Transactions on Algorithms* 13(1): Article #17, 2016.

- [19] D. Harris and A. Srinivasan. A constructive Lovász local lemma for permutations, *Theory of Computing* 13(17): 1-41, 2017.
- [20] D. Harris and A. Srinivasan. Algorithmic and enumerative aspects of the Moser-Tardos distribution, *ACM Transactions on Algorithms (TALG)* 13(3): Article #33, 2017.
- [21] P. Haxell. A condition for matchability in hypergraphs, *Graphs and Combinatorics* 11(3): 245-248, 1995.
- [22] P. Haxell. A note on vertex list colouring, *Combinatorics, Probability, and Computing* 10(4): 345-347, 2001.
- [23] P. Haxell. On the strong chromatic number, *Combinatorics, Probability, and Computing* 13(6): 857-865, 2004.
- [24] P. Haxell. An improved bound for the strong chromatic number, *Journal of Graph Theory* 58(2): 148-158, 2008.
- [25] P. Haxell. On forming committees, *The American Mathematical Monthly* 118(9): 777-788, 2011.
- [26] A. Lo and N. Sanhueza-Matamala. An asymptotic bound for the strong chromatic number, *Combinatorics, Probability, and Computing*. 28(5): 768-776, 2019.
- [27] P. Loh and B. Sudakov. On the strong chromatic number of random graphs, *Combinatorics, Probability, and Computing*. 17(2): 271-286, 2007.
- [28] M. Molloy and B. Reed. The strong chromatic number, In *Graph Colouring and the Probabilistic Method*, Springer, Berlin, 2002.
- [29] R. Moser and G. Tardos. A constructive proof of the general Lovász Local Lemma, *Journal of the ACM*, 57(2): Article #11, 2010.
- [30] L. Rabern. Coloring graphs with dense neighborhoods, *Journal of Graph Theory* 76(4), 323-340, 2014.
- [31] T. Szabó and G. Tardos. Extremal problems for transversals in graphs with bounded degree, *Combinatorica*. 26(3): 333-351, 2006.