Electronic Colloquium on Computational Complexity, Report No. 138 (2011)

Complexity Lower Bounds through Balanced Graph Properties

Guy Moshkovitz*

Abstract

In this paper we present a combinatorial approach for proving complexity lower bounds. We mainly focus on the following instantiation of it. Consider a pair of properties of m-edge regular hypergraphs. Suppose they are "indistinguishable" with respect to hypergraphs with $m - t$ edges, in the sense that every such hypergraph has the same number of super-hypergraphs satisfying each property. Roughly speaking, we show that finding a pair of distinct such properties implies an $m/(t-1)$ lower bound on the rank of explicit tensors.

We also show, albeit non-explicitly, that near-optimal rank lower bounds can be obtained in this manner. Furthermore, we consider the $t = 2$ case and prove that it already implies non-trivial lower bounds. In particular, we derive a (tight) lower bound of $3n/2$ on the rank of $n \times n \times n$ tensors naturally associated with hypergraph forests (which apparently was not known before; in fact, our bound also applies to the so-called *border* rank, and as such, is not far from the best lower bounds known).

1 Introduction

It is a fundamental and long-standing open challenge to find a polynomial that provably has no small *arithmetic circuit*. Roughly speaking, an arithmetic circuit is a description of a polynomial using arithmetic operations—addition, subtraction, and multiplication—by applying those repeatedly, starting from the variables (and any constant from the underlying field). Despite considerable efforts, the number of arithmetic operations required to thus compute a polynomial—its *complexity*—remains poorly understood. Specifically, the best lower bound known on the complexity of an *explicit*[1] n-variate polynomial is slightly super-linear $\Omega(n \log n)$, even though an easy counting argument shows that the complexity of almost all n-variate polynomials of degree $O(n)$ is exponential $2^{\Omega(n)}$. In fact, relatively little is known even on the more restricted *arithmetic formula*, in which a previously computed polynomial can be used as an input to only one subsequent arithmetic operation.

In this paper, we define and study a notion of "balance" of hypergraph properties, and show that it has tight connections with complexity lower bounds. The exact meaning of "balance" depends on the model of computation for which we would like to prove lower bounds, be it

*Blavatnik School of Computer Science, Tel Aviv University, Israel. Supported by the Israel Science Foundation.

[1]A polynomial p (or more accurately, a sequence $(p_n)_{n \in \mathbb{N}}$ of n-variate polynomials) is considered explicit if, roughly, it is of degree polynomial in n, and there is a polynomial time Turing machine that when specified a monomial, outputs its coefficient in p.

ISSN 1433-8092

an arithmetic circuit, an arithmetic formula, or other natural models. For concreteness and simplicity we shall focus in this introduction, and throughout most of the paper, on the definition of "balance" that implies lower bounds on *tensor rank*. (We stress that almost every result we present in this paper is applicable, usually verbatim, to other models of computation; see Section 7 for a discussion.)

Tensors. Let $\mathbb{F}$ denote an arbitrary field, and let $A = (a_e)_{e\in[n]^d}$ be a d-tensor with entries in $\mathbb{F}$. We say that a nonzero A is of rank 1 if there are d vectors $v^1, \ldots, v^d \in \mathbb{F}^n$ such that $A = v^1 \otimes \cdots \otimes v^d$ (where $\otimes$ denotes tensor product, i.e., $A = (v^1_{e_1} \cdots v^d_{e_d})_{e\in[n]^d}$). The *rank* of A is the least r so that A is a sum of r rank-1 tensors.[2] Although the rank of matrices, or 2-tensors, is well understood, not much is known on the rank of d-tensors for $d > 2$. For instance, while it is known that almost all d-tensors are of rank $\approx n^{d-1}/d$ (note that the maximum possible rank is clearly at most n^{d-1}), no rank lower bound better than (a folklore) $\Omega(n^{\lfloor d/2\rfloor})$ is known for an explicit tensor. In particular, no explicit 3-tensor of super-linear $\omega(n)$ rank is known.

Perhaps surprisingly, rank lower bounds imply in some cases lower bounds for general arithmetic computations. Results of Strassen [26] and Baur and Strassen [4] show that a lower bound on the rank of a 3-tensor implies the same lower bound, up to a constant factor, on the complexity of computing the naturally-associated polynomial by an arithmetic circuit (this has the potential to lead to circuit lower bounds of up to $\Omega(n^2)$). For d-tensors, a result of Raz [22] shows that a rank lower bound of n^k implies a lower bound of $n^{\Omega(\log(\frac{d}{d-k}))}$ on the complexity of computing the associated polynomial by an arithmetic *formula*, assuming $d \le \log n/\log\log n$ (this has the potential to lead to super-polynomial formula lower bounds). Currently, no known rank lower bound is strong enough for these results to imply interesting general arithmetic lower bounds.

1.1 Our results

A hypergraph[3] is *d-partite* if its vertex set is partitioned into d classes so that every edge includes exactly one vertex from each class (e.g., 2-partite hypergraphs are the usual bipartite graphs). The *adjacency tensor* of a d-partite hypergraph, generalizing the *adjacency matrix* of a bipartite graph, has the value 1 in the entries corresponding to the edges, and 0 elsewhere. The *degree list* of a hypergraph H, denoted $\mathfrak{L}(H)$, is the *labeled* sequence specifying the degree[4] of each vertex in H. (Clearly, different hypergraphs may have the same degree list.)

An important concept in this paper is that of a *t-scope*, which is a set of d-partite hypergraphs of the form[5]

$$\{\, G_0 \mathbin{\dot\cup} H \mid \mathfrak{L}(H) = L \,\}$$

[2] This notion of tensor computation can be viewed as a restriction of an arithmetic formula, and in that context is sometimes referred to as a depth-3 set-multilinear formula; see, e.g., [22] for a discussion.

[3] A (finite, simple) hypergraph is a collection of subsets, called edges, of some finite set of vertices. Unless otherwise specified, we implicitly assume that all hypergraphs considered are over the same vertex set.

[4] The degree of a vertex is the number of edges containing it.

[5] The symbol $\dot\cup$ stands for disjoint union; i.e., $G_0 \cap H = \emptyset$.

for a d-partite hypergraph G_0 and a degree list L of t-edge d-partite hypergraphs (in other words, it is the set of super-hypergraphs of G_0 with a given degree list, namely, $\mathfrak{L}(G_0) + L$). As a trivial example, note that a degree list L of a 1-edge hypergraph clearly determines a single hypergraph H. Hence, any 1-scope contains either a single member, or none (in case G_0 already has that edge). A more interesting case is that of 2-scopes; for example, if a degree list L has two vertices of degree 1 from each vertex class (and all others of degree 0), then there are exactly 2^{d-1} hypergraphs H (having two edges) such that $\mathfrak{L}(H) = L$. Hence, any 2-scope with such an L contains at most 2^{d-1} hypergraphs, and in fact it is easy to see that the same holds for any 2-scope.

Let $\mathcal{P}, \mathcal{P}'$ be two disjoint families (or properties) of d-partite hypergraphs over a given vertex set. Suppose that for every t-scope, the number of hypergraphs from $\mathcal{P}$ that are in the scope is the same as the number of hypergraphs from $\mathcal{P}'$ that are in the scope. In such a case, we refer to the union $\mathcal{P} \cup \mathcal{P}'$ as a *t-balanced* family. We say that a hypergraph is *t-balanceable* if there is a t-balanced family containing it. Our main result is the following (see Theorem 3.9).

Theorem 1.1 (Main Theorem)**.** *The rank of the adjacency tensor of any t-balanceable m-edge hypergraph is at least $\left\lceil \frac{m}{t-1} \right\rceil$, over any field.*[6]

Very intuitively, the above bound follows as no rank-1 tensor is able to "capture" any t of the m edges in the balanceable hypergraph.

It should be noted that finding a balanced family actually implies lower bounds for many more tensors, as follows (see Theorem 3.11).

Theorem 1.2. *Let $\mathcal{P} \cup \mathcal{P}'$ be a t-balanced family, and suppose that its members have m edges. Then the $\left\lceil \frac{m}{t-1} \right\rceil$ rank lower bound holds for the adjacency tensor of any d-partite hypergraph whose number of subgraphs from $\mathcal{P}$ and from $\mathcal{P}'$ differ, assuming the difference is not a multiple of the characteristic of the field.*

Weak balance. It is also possible to relax our definition and deduce lower bounds over fields of characteristic 2. We say that a family of d-partite hypergraphs is *weakly* t-balanced if the number of its members in every t-scope is even.[7] Notice that any t-balanced family is clearly also weakly t-balanced. Naturally, we say that a hypergraph is *weakly t-balanceable* if there is a weakly t-balanced family containing it. This definition is sufficient to deduce the following lower bound (see Theorem 3.10).

Theorem 1.3 (Main Theorem II)**.** *The rank of the adjacency tensor of any weakly t-balanceable m-edge hypergraph is at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field of characteristic 2.*

Observe that a lower bound over the binary field GF(2) also implies a lower bound over the integers: if a tensor A with $\{0, 1\}$ entries is of rank at most r over the integers, meaning it can

[6]Note that this bound is well defined as there are no 1-balanceable hypergraphs, which follows from the fact that every hypergraph appears in a singleton 1-scope, and thus cannot appear in only one of $\mathcal{P}$ or $\mathcal{P}'$.

[7]This definition is actually a bit simplified; the actual definition involves considering s-scopes for all $s \geq t$ (see Definition 3.7), but as we will show in Corollary 4.12, for t a power of two, it turns out to be equivalent to the simpler definition given here.

be represented as a sum of r rank-1 tensors with integer entries, then reducing the latter modulo 2 implies that A is of rank at most r over GF(2) as well.

When constructing balanceable hypergraphs in this paper, we begin by first showing that our hypergraphs are weakly balanceable—which already implies rank lower bounds—and only then try to find a good bipartition of the containing family (i.e., into a pair $\mathcal{P}, \mathcal{P}'$ as above), thereby extending the lower bound result so as to hold over any field.

Constructing balanceable hypergraphs. One might wonder how large can lower bounds deduced using the above theorems be. We show, albeit in a non-explicit manner, that there exist balanceable hypergraphs with parameters that imply almost optimal rank lower bounds (see Theorem 3.12).

Theorem 1.4. *For every $d \geq 2$, and every sufficiently large n, there exist t-balanceable d-partite hypergraphs with m edges and n vertices in each vertex class, for some m and t such that $\left\lceil \frac{m}{t-1} \right\rceil = \Omega(\frac{n^{d-1}}{d \log d})$.*

Of course, our goal is to find *explicit* lower bounds. To that end, we construct some balanceable hypergraphs in Section 5; in particular, we construct 2-balanceable hypergraphs. Observe that showing a hypergraph to be 2-balanceable exactly determines the rank of its adjacency tensor, since the adjacency tensor of a hypergraph with m edges is clearly of rank at most m (in other words, showing a hypergraph to be 2-balanceable implies that its adjacency tensor is "full rank"). We prove, for example, that any d-partite forest (see Subsection 5.2 for definitions) is weakly 2-balanceable, assuming $d > 2$. A simple argument shows that any such hypergraph has up to $(dn-1)/(d-1) \leq (3n-1)/2$ edges, where n is the number of vertices in a vertex class. We also show that any *binary* 3-partite tree, which is a connected forest of maximum degree 2, is (non-weakly) 2-balanceable. Since such hypergraphs having $(3n-1)/2$ edges exist, we deduce a tight rank lower bound of $(3n-1)/2$ on their adjacency tensors that holds over any field. In addition, we show that these constructions cannot be improved by much, as any (weakly) 2-balanceable 3-partite hypergraph has in fact at most $3n$ edges, and hence in order to obtain lower bounds greater than $3n$ for 3-tensors, one must consider t-balanceable hypergraphs for $t \geq 3$.

Known results on tensor rank. Previous work on the rank of tensors mainly focused on some 3-tensors of interest. We next describe a particularly interesting one. Consider the problem of computing the product of two $\ell \times \ell$ matrices $(a_{ij})_{i,j}$ and $(b_{ij})_{i,j}$, that is, the $\ell \times \ell$ matrix $(\sum_{k=1}^{\ell} a_{i,k} b_{k,j})_{i,j}$. The arithmetic complexity of this problem has been studied extensively over the years; it is known, following the above-mention [26, 4], to equal, up to a constant factor, to the rank of the naturally-associated size-ℓ^2 3-tensor. The best lower bounds known on the rank of this tensor are $2.5\ell^2 - \ell$ over any field, shown by Bläser [6] (whose proof relied on showing that in any basis of the space of matrices, there must be a small subset that spans matrices satisfying a certain invertibility condition), and $3\ell^2 - O(\ell^{\frac{5}{3}})$ over GF(2), shown by Shpilka [25] (combining the above methods with techniques from linear codes). On the other hand, the best rank upper bound known is $O(\ell^{2.38})$ by Coppersmith and Winograd [8].

There are also results on the related notion of *border rank* of a tensor. Over the complex numbers, the border rank of a tensor A is the smallest r so that every polynomial, acting on tensors, that evaluates to zero on all tensors of rank at most r, also vanishes on A.[8] It is clear that the border rank of a tensor is never larger than its rank. The best lower bound known on the border rank of the matrix multiplication tensor is $3\ell^2/2+\ell/2-1$; this was shown by Lickteig [16] (the proof is based on an upper bound on the dimension of a related linear subspace).

As for tensors other than that of matrix multiplication, the best rank lower bound known on $n \times n \times n$ tensors, shown in [3], is $3n - O(\log n)$, and it holds over any field. For lower bound proofs that apply to a large family of (explicit) 3-tensors, see [11, 13] where rank lower bounds of $3n/2$ are shown for $2 \times n \times n$ tensors (using linear-algebraic tools). Furthermore, Strassen [28] gave $3n/2$ *border* rank lower bounds for $3 \times n \times n$ tensors; this was later improved by Griesser [10], but as far as we know, no explicit border rank lower bound better than $2n$ for $n \times n \times n$ tensors is known. We mention that although we obtain a modest lower bound of $3n/2$ on the (in fact, border-) rank of tensors, our proof applies to tensors corresponding to hypergraph forests, which seems more natural compared to previous lower bounds.

Techniques. To prove Theorems 1.1, 1.2, and 1.3 we use the following algebraic zeroing argument. Consider the rd vectors that specify a rank-r computation of a d-tensor. It is not hard to see that the functions mapping a rank-r computation—viewed as a sequence of rdn field elements—to each entry of the computed tensor are in fact polynomials. If we are able to construct a polynomial Q acting on d-tensors (i.e., Q has n^d variables) such that composing it on the aforementioned n^d polynomials yields the zero polynomial, then rank lower bounds follow—any tensor on which Q does not vanish must be of rank greater than r. We show that by using our notion of balancedness, we can construct such a polynomial Q, and which is also multilinear and homogeneous. Using the latter properties of Q, non-zeros can be easily found. For example, using the fact that Q is multilinear, it is not hard to see that any of its monomials naturally corresponds to a subset of the entries of a tensor; hence a tensor whose nonzero entires are given by such a subset must, by homogeneity, be a non-zero of Q.

Polynomials that vanish on low-complexity arithmetic computations were considered in the past, e.g., by Motzkin [18], Belage [5], Strassen [27], and Lipton [17], who proved not-quite-explicit complexity lower bounds for univariate polynomials. For instance, Strassen gave a near-optimal lower bound on univariate polynomials whose sequence of (integer) coefficients grows double-exponentially. This result was proved by observing that there exists a polynomial H that vanishes on sequences of coefficients of low-complexity polynomials such that H is of moderate degree and has small integer coefficients. It was then shown that such an H cannot vanish when evaluated on double-exponentially-sized inputs, which completed the proof. (In fact, Theorem 3.12, showing the existence of a balanced family with good parameters, implies a somewhat stronger existence results for such an H.) Also, Strassen's result [28] mentioned above was proved by explicitly describing such a polynomial, vanishing on low-rank $3 \times n \times n$ tensors.

[8] It can also be defined as the least r so that the tensor is the limit of a sequence of tensors of rank at most r. Thus, intuitively, a border rank lower bound on a tensor says that it cannot even be "approximated" by low-rank tensors. For definitions over other fields see, e.g., [15].

In another, more recent work, by Raz [21], such polynomials were used to show that a certain polynomial mapping "eludes" being computed by small bounded-depth arithmetic circuits.

The polynomials we consider here are, compared to those used in the aforementioned papers, more "well behaved" in that they are multilinear, homogeneous, and with unity coefficients. Our results are based on the fact that the action of such well-behaved polynomials can be described using t-scopes, for appropriate values of t.

Let us mention that some approaches for proving complexity lower bounds were suggested in the past using, e.g., Raz's elusive functions in [21], Valiant's rigid matrices in [29], as well as the algebraic geometry approach in [20]. The approach suggested in this paper is purely (hyper-) graph-theoretical, and seems amenable to tools and techniques from that area.

As for the techniques we use in our constructions of balanceable hypergraphs, they are based on the following observation. Think of hypergraphs from the same scope as being obtained from one another by applying a "rewiring operation", replacing a subgraph consisting of t edges with a different subgraph having the same degree list. To prove that any d-partite forest ($d > 2$) is weakly 2-balanceable, we show that every two of its edges can be rewired so as to obtain a different forest. More formally, we give, for every 2-scope, a fixed-point-free involution on the forests contained there.

While we do not know if all these forests are in fact 2-balanceable, we are able to show this—as mention above—for binary trees. To do so, we define a notion of parity for binary trees, and prove that applying the aforementioned rewiring involution always changes the parity. This implies that the families of binary trees of even and of odd parity form a 2-balanced family. This implies that each of their members is 2-balanceable, as desired.

1.2 Open questions

This paper raises some open questions of interest.

- The main open question is whether we can find *explicit* t-balanceable hypergraphs whose number of edges m is large relative to t. In particular, can we find such a hypergraph for which m/t is super-linear in the number of vertices? Showing this for 3-partite hypergraphs would lead to the first super-linear arithmetic circuit complexity lower bound for polynomials of constant degree.

- A more modest goal would be to find explicit t-balanceable 3-partite hypergraphs for which $m/t > 3n$, where n is the number of vertices in a vertex class (note that this is not possible for $t = 2$ by Corollary 5.14). This would improve the best rank lower bounds known.

- Can we use our approach to prove new lower bounds for specific tensors/polynomials of interest, such as the matrix multiplication tensor?

- Does the "natural proof" barrier pose significant restrictions on possible constructions of balanceable hypergraphs? Razborov and Rudich [23] show that, under a reasonable conjecture, any family of functions of high *Boolean* circuit complexity is, roughly speaking, small (relative to the number of functions) or "complicated" (deciding membership in it is computationally hard). Although this barrier is currently not known to apply to lower bounds on arithmetic complexity (or restrictions thereof), it might still have interesting implications for families of balanceable hypergraphs.

- Another interesting question is whether the proof of Theorem 5.10, showing that *binary* d-partite trees ($d > 2$) are 2-balanceable, can be extended to all trees/forests.

1.3 Outline

The rest of the paper is organized as follows. After some preliminaries, in which we define the complexity of a partite hypergraph as the rank of its adjacency tensor, we formally define rewirings and balanced families of hypergraphs in Section 3. We state their connection to complexity lower bounds, deferring the proof to Section 6 (where we also discuss some extensions); moreover, we prove the existence of balanced families with good parameters. In Section 4 we present some properties of balanced families of hypergraphs, as well as give an equivalent definition in terms of local hypergraph operations. We present explicit examples of balanced families of hypergraphs in Section 5. Results for additional computational models are discussed in Section 7.

2 Preliminaries

For the sake of brevity, and unless otherwise specified, we always take *graph* (sometimes *partite graph*) to mean a d-partite hypergraph, for some $d \geq 2$.

We say that a partite graph is *n-bounded* if there are n vertices in each vertex class. The size (or cardinality) of a graph G is, by definition, its number of edges, denoted $|G|$. A *t-subgraph* is a subgraph of cardinality t.

For a positive integer n, we denote $[n]$ the set $\{1, 2, \ldots, n\}$. We denote $\mathbb{N}$ the set of non-negative integers, and GF(2) the binary field. We use $\log(\cdot)$ to denote logarithm in base 2. We define the *complexity* of a partite graph as the rank of its adjacency tensor. It is easy to verify the following nice property: The complexity of any induced subgraph[9] of a graph G is at most the complexity of G. Finally, we record the easy fact that the complexity of any n-bounded d-partite graph is at most n^{d-1}.

3 Basics

This section contains the basic definitions we use, namely, rewirings of partite graphs and balanced families of graphs. We state the connection between these definitions and tensor rank lower bounds, as well as prove that balanced families of graphs with good parameters exist.

[9] An induced subgraph of a (hyper-) graph is obtained by deleting a subset of the vertices (formally, by removing every edge containing a vertex from the subset).

3.1 Rewirings of partite graphs

Recall that the *degree list* of a hypergraph H, denoted $\mathfrak{L}(H)$, is the labeled sequence specifying the degree of each vertex in H. We first introduce the following straightforward terminology.

Definition 3.1. Two d-partite graphs are *comparable* if they have the same degree list.

That is, two graphs are comparable if they are spanned by the same vertices, and moreover, the degree of each vertex in one graph is equal to its degree in the other graph. Observe that comparable graphs have the same cardinality (i.e., same number of edges). Let us also mention that for bipartite graphs, the equivalence classes of the comparable equivalence relation were studied extensively over the years, under many guises, usually referred to as *classes of contingency tables* or *classes of* $(0,1)$*-matrices with fixed row and column sums*; see, e.g., [7] for a survey.

Given a partite graph, we define a *rewiring* of a subgraph of it to be any graph obtained by replacing the subgraph with a comparable one.

Definition 3.2 (Rewiring)**.** Let G be a partite graph and let H be a subgraph of G. A *rewiring* of H in G is any graph of the form $(G \setminus H) \cupdot H'$ where H' is comparable to H.

Thus, intuitively, a rewiring of a subgraph H in G is obtained from G by "rearranging" the vertices of H between the edges, while preserving the degree of each vertex. Clearly, G itself is also a rewiring of any H in G; unless H contains at most one edge, there are in general many other rewirings. We emphasize that any rewiring of a subgraph in G is, in particular, comparable to G.

Remark. Not every graph H' that is comparable to H has a corresponding rewiring of H in G. This is precisely because H' might include edges from $G \setminus H$, so that replacing H with H' in G would result in multiple edges.

We next consider an especially important example, showing that complete subgraphs (i.e., cliques) have no non-trivial rewiring. A *complete d-partite graph* $K(t_1, \ldots, t_d)$—whose ith vertex class has t_i vertices—contains all possible edges, where every edge includes one vertex from each vertex class (equivalently, this is the Cartesian product of the d vertex classes).

Claim 3.3. *Let H be a complete partite graph, and suppose that H is contained in a partite graph G. Then there is no rewiring of H in G other than G.*

Proof. We prove that any complete graph is comparable only to itself. Observe that the degree of each vertex in a complete graph is the maximum possible among all partite graphs on the same vertex set. Thus, for any vertex in a graph comparable to a complete graph, every possible edge that includes it must appear. This proves the claim. $\square$

Let us introduce another useful terminology, as follows.

Definition 3.4 (Scope)**.** For a graph G and a subgraph $H \subseteq G$ of cardinality t, we refer to the set of all rewirings of H in G as a *t-scope*. If H is of cardinality at least t, we use *t^+-scope*.

Note that every t-scope Ψ can be parameterized by a pair (G_0, L), where G_0 is a graph and L is a degree list of size-t graphs, so that

$$\{\, G_0 \mathbin{\dot\cup} H \mid \mathfrak{L}(H) = L \,\}.$$

We sometimes refer to G_0 as the "fixed part" of the scope. Let us record the following basic property of scopes.

Fact 3.5. *All graphs in the same scope are comparable, and in particular, of the same cardinality.*

As all graphs in a scope Ψ belong to the same equivalence class $\mathcal{C}$ of the comparable relation (containing all graphs with a given degree list), we may in fact characterize Ψ as the subset of $\mathcal{C}$ obtained by "fixing" a subgraph G_0. That is, Ψ is the set of supergraphs of G_0 in $\mathcal{C}$,

$$\Psi = \{\, G \in \mathcal{C} \mid G_0 \subseteq G \,\}.$$

Also—avoiding any mention of G_0 or L—notice that for two graphs G and G' of cardinality at least t, there is a t-scope containing both if and only if they are comparable and their difference is of cardinality at most t (i.e., $|G \setminus G'| = |G' \setminus G| \leq t$).

3.2 Balanced families of graphs

Our main objects of study are defined as follows.

Definition 3.6. A family of graphs $\mathcal{A}$ is *t-balanced* if it can be bipartitioned $\mathcal{A} = \mathcal{A}^+ \mathbin{\dot\cup} \mathcal{A}^-$ so there is an equal number of graphs from $\mathcal{A}^+$ and from $\mathcal{A}^-$ in every t^+-scope.

That is, a family of graphs—henceforth abbreviated *family*—is balanced if it consists of two disjoint parts that are "indistinguishable" when intersected with every t^+-scope. As a consequence of Fact 3.5, any t-balanced family can be decomposed into disjoint t-balanced families, each containing graphs with a different degree list. In that sense, it is mainly interesting to consider balanced families that contain only graphs with the same degree list (e.g., k-regular graphs for any k), and in particular of the same cardinality. Notice that if $\mathcal{A}$ is a family of comparable size-m graphs, then $\mathcal{A}$ is t-balanced if and only if for every graph G_0 of cardinality at most $m - t$, the number of supergraphs of G_0 that lie in $\mathcal{A}^+$ (i.e., members of $\mathcal{A}$ that contain G_0) is the same as the number of supergraphs of G_0 that lie in $\mathcal{A}^-$.

Also interesting is the following weaker definition.

Definition 3.7. A family of graphs is *weakly* t-balanced if the number of graphs from the family in every t^+-scope is even.

For equivalent definitions, see Section 4 (where we show, e.g., that one may consider only t-scopes, as opposed to t^+-scopes, in the definition of a t-balanced family).

There are clearly no nonempty 1-balanced families—not even in the weaker sense—as any 1-scope contains (at most) a single graph. On the other extreme, m-balanced families containing size-m graphs are abound; they are easily characterized as follows.

Example 3.8. Any graph G of cardinality m appears in a single m-scope, which contains every graph comparable to G. Therefore, any two comparable graphs of cardinality m form an m-balanced family. More generally, a family $\mathcal{A}$ of size-m graphs is m-balanced if and only if the number of graphs from $\mathcal{A}$ with any given degree list is even.

In Claim 5.3 we give a simple example of a (rather large) t-balanced family of size-m graphs where t is in fact somewhat smaller than m. On the negative side, following Claim 3.3, no graph that contains a complete subgraph of cardinality at least t is a member of a (weakly) t-balanced family. These examples naturally raise the question: can one find a t-balanced family containing large graphs, for t that is small relative to the size of the graphs? The following theorem shows that a positive answer to this question implies large complexity lower bounds. Call a graph t-*balanceable* if it is a member of a t-balanced family.

Theorem 3.9 (Main Theorem)**.** *The complexity of any t-balanceable graph of cardinality m is at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field.*

One can very intuitively interpret the above theorem as follows. Notice that when expressing the adjacency tensor of a size-m graph as a sum of less than $\left\lceil \frac{m}{t-1} \right\rceil$ rank-1 tensors, there must be a rank-1 tensor that in a sense "captures" at least t edges. Roughly speaking, a t-balanceable graph has no t edges that can be thus captured. In other words, the distinction between the two families $\mathcal{A}^+$ and $\mathcal{A}^-$ is "too subtle" for a low-rank computation to make (and for combinatorial reasons, i.e., this holds regardless of the field underlying the computation).

In fact, lower bounds also follow from the weaker notion of balance. Call a graph *weakly* t-balanceable if it is a member of a weakly t-balanced family.

Theorem 3.10 (Main Theorem II)**.** *The complexity of any weakly t-balanceable graph of cardinality m is at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field of characteristic* 2*.*

We defer the proofs of the above theorems to Section 6.

Remark. It follows that any graph G that is balanceable (in a non-trivial manner, i.e., t-balanceable for some $t \leq |G|$) is in particular of complexity at least 2, since in this case $|G| > t-1$, which implies $\left\lceil \frac{|G|}{t-1} \right\rceil \geq 2$. Furthermore, complete (partite) graphs are, by definition, of complexity exactly 1. By the proof of Claim 3.3, any complete graph is comparable only to itself, which means that no complete graph is (non-trivially) balanceable, as expected.

Further implications of balanced families. As discussed above, showing that a family of graphs is balanced implies complexity lower bounds for each of its members. In fact, the same lower bound applies also to any graph containing as an *induced* subgraph a graph in the family (since the complexity of a graph is at least the complexity of any of its induced subgraphs). The following theorem shows another way to derive lower bounds on the complexity of larger graphs from a balanced family. To illustrate its usefulness, consider a graph that contains as a non-induced subgraph a single t-balanceable graph, as well as a complete subgraph of cardinality at least t. Notice that such a graph cannot be t-balanceable—even in the weak sense—and hence

Theorems 3.9 and 3.10 cannot be used to deduce lower bounds on its complexity. Nevertheless, as we now show, the same complexity lower bound still applies to it.

Theorem 3.11. *Let $\mathcal{A}$ be a weakly t-balanced family of size-m graphs. Any graph whose number of subgraphs from $\mathcal{A}$ is odd is of complexity at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field of characteristic 2. Moreover, if $\mathcal{A}$ is t-balanced then the same holds also over any field of characteristic zero. In fact, denoting $\mathcal{A} = \mathcal{A}^+ \cup \mathcal{A}^-$ the associated bipartition, the lower bound holds already for graphs whose number of subgraphs from $\mathcal{A}^+$ and from $\mathcal{A}^-$ differ, assuming the difference is not a multiple of the characteristic of the field.*

As a corollary, it clearly follows that a graph containing an induced subgraph of cardinality k that satisfies the conditions in Theorem 3.11 is of complexity at least $\left\lceil \frac{k}{t-1} \right\rceil$.

3.3 Existence of balanced families

We proceed to show that nonempty t-balanced families exist with t small relative to the size of their graphs.

Theorem 3.12. *For every $d \geq 2$, and every sufficiently large n, there exist t-balanceable n-bounded d-partite graphs of cardinality m for some $m = \Omega(n^d)$ and $t = O(nd \log d)$.*

Notice that balanceable graphs with parameters as above imply, using Theorem 3.9, a rank lower bound of $\frac{m}{t} = \Omega(\frac{n^{d-1}}{d \log d})$, over any field. Since the complexity of any n-bounded d-partite graph is at most n^{d-1}, we conclude that almost optimal complexity lower bounds can be obtained in this manner.

Proof. Let $t \leq m$ be positive integers. Let $\mathbf{G}$ denote the set of all n-bounded d-partite graphs of cardinality m (over a fixed vertex set), and let $\mathbf{S}$ denote the set of all t-scopes containing such graphs. Define the function $\Lambda : 2^{\mathbf{G}} \to \mathbb{N}^{\mathbf{S}}$, mapping families of graphs to sequences of nonnegative integers, by

$$\Lambda(\mathcal{F}) = \big(|\mathcal{F} \cap \Psi|\big)_{\Psi \in \mathbf{S}}.$$

Suppose there are two distinct families $\mathcal{F}_1, \mathcal{F}_2 \in \mathbf{G}$ that are mapped by Λ to the same sequence. In Proposition 4.7 we show the easy fact that for a pair of families containing the same number of members in every t-scope, their symmetric difference is a t-balanced family. Since $\mathcal{F}_1 \neq \mathcal{F}_2$, the symmetric difference $\mathcal{F}_1 \triangle \mathcal{F}_2$ is nonempty. Therefore, to complete the proof we need to show the existence of such a pair of families.

Let M denote the cardinality of the largest t-scope (containing size-m n-bounded d-partite graphs). Notice that the image of Λ is of cardinality at most $M^{|\mathbf{S}|}$. Using the pigeonhole principle, it follows from the above discussion that if $M^{|\mathbf{S}|} < 2^{|\mathbf{G}|}$ then there exists a nonempty t-balanced family. We therefore need to show that $|\mathbf{S}| \cdot \log M < |\mathbf{G}|$.

Let us estimate the above quantities. Recall that the t-scopes containing size-m graphs are in bijection with the pairs (G_0, L), where G_0 is a graph of size $m - t$, and L is a degree list of size-t graphs. Since the underlying vertex set has n vertices in each of the d vertex classes, it follows that the number of such degree lists is at most $\binom{n+t-1}{t}^d$ (since any degree list specifies, for each vertex class, a size-t multiset of vertices). Therefore, $|\mathbf{S}| \leq \binom{n^d}{m-t}\binom{n+t-1}{t}^d$. Also, we clearly have that $|\mathbf{G}| = \binom{n^d}{m}$. Finally, since any t-scope contains only rewirings of size-t graphs, it follows that $M \le \binom{n^d}{t}$, being the number of size-t n-bounded d-partite graphs.

Before continuing, we will need two basic bounds. First, we give an easy lower bound on the ratio of two "close" binomial coefficients.

Claim 3.13. *Let c, k, n be positive integers and suppose that $c \le k \le n/3$. Then*

$$\frac{\binom{n}{k}}{\binom{n}{k-c}} > 2^c.$$

Proof. For all $1 \le i \le n/3$,

$$\binom{n}{i} = \frac{n-i+1}{i}\binom{n}{i-1} > 2\binom{n}{i-1}.$$

□

Next is an upper bound on the number of large multisets over a given ground set.

Claim 3.14. *For all positive integers $n \le k$,*

$$\binom{n+k-1}{k} < \left(6\frac{k}{n}\right)^n.$$

Proof. Using the bound $\binom{a}{b} \le (e\frac{a}{b})^b$, we have

$$\binom{n+k-1}{k} \le \binom{n+k}{k} = \binom{n+k}{n} \le \left(e\left(1+\frac{k}{n}\right)\right)^n \le \left(2e\frac{k}{n}\right)^n.$$

□

We are now ready to prove that $|\mathbf{S}| \cdot \log M < |\mathbf{G}|$; as discussed above, it suffices to show that

$$\binom{n^d}{m-t}\binom{n+t-1}{t}^d \cdot \log\binom{n^d}{t} < \binom{n^d}{m}.$$

Assume $t \le m \le n^d/3$. By Claim 3.13,

$$\frac{\binom{n^d}{m}}{\binom{n^d}{m-t}} > 2^t.$$

Put $t = ndq$ where q is a positive integer. By Claim 3.14,

$$\binom{n+t-1}{t} < \left(6\frac{t}{n}\right)^n = (6dq)^n.$$

Moreover,

$$\log\binom{n^d}{t} \le t\log\left(e\frac{n^d}{t}\right) \le t\log(n^d) \le t^2 = n^2(dq)^2 \le 6^n(dq)^2 \le (6dq)^n.$$

Therefore, it is enough to show that $(6dq)^{nd+n} \le (6dq)^{2nd}$ is less than $2^t = 2^{ndq}$, or equivalently, $6dq < 2^{q/2}$. It is easy to check that choosing, e.g., $q = \lceil 16\log d\rceil$ (that is, $t = nd \cdot \lceil 16\log d\rceil$) and $m = n^d/3$ completes the proof. □

4 Properties and equivalent definitions

In this section we present some basic properties of balanced families, as well as some equivalent definitions. For explicit examples of balanced families, the reader may skip directly to Section 5. We begin with a simple bound applicable to balanced families that contain large graphs. Note that it follows from Claim 3.3 that no (weakly) t-balanceable graph contains a complete subgraph of cardinality at least t. In particular, any graph isomorphic to $K(1,\ldots,1,t,1,\ldots,1)$ for some $t \geq 1$, which we refer to as a *t-cluster*, must be excluded from any (weakly) t-balanceable graph. We first need the following easy fact concerning the extremal number of clusters.

Lemma 4.1. *Every n-bounded d-partite graph of cardinality larger than $(t-1)n^{d-1}$ contains a t-cluster.*

Proof. Given a graph, define the degree of a set of vertices—each from a different vertex class—as the number of edges that contain all of them. For any size-m graph and for every k vertex classes, there must be a set of vertices, one from each of the k vertex classes, whose degree is at least $\lceil \frac{m}{n^k} \rceil$, since $\frac{m}{n^k}$ is the average degree of such a set. In particular, for $k = d-1$, there must be at least $\lceil \frac{m}{n^{d-1}} \rceil$ edges whose common intersection is of size at least (and thus exactly) $d-1$. Since the latter is clearly equivalent to a cluster, any graph of cardinality $m > (t-1)n^{d-1}$ must contain a t-cluster, as required. □

Corollary 4.2. *Any (weakly) t-balanceable n-bounded d-partite graph is of cardinality at most $(t-1)n^{d-1}$.*

We remark that one can alternatively prove the result above as a corollary of Theorem 3.9, by using the fact that the complexity of any n-bounded d-partite graph is at most n^{d-1}.

More generally, the extremal number of any type of complete graph $K(n_1,\ldots,n_d)$ where $\prod_i n_i \geq t$ (i.e., not just t-clusters) is an upper bound on the cardinality of a (weakly) t-balanceable graph. The extremal numbers of complete partite graphs have been investigated extensively (e.g., for bipartite graphs, finding these extremal numbers is known as the Zarankiewicz problem). See, for instance, the paper [9] by Erdős for some results.

Symmetric differences. We now show that, in the definition of a balanced family, it is not necessary to require that its subfamilies $\mathcal{A}^+$ and $\mathcal{A}^-$ are disjoint. In other words, any pair of not necessarily disjoint "indistinguishable" families form a balanced family.

Proposition 4.3. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two (not necessarily disjoint) families of graphs. Suppose that for every t^+-scope Ψ it holds that $|\Psi \cap \mathcal{A}_1| = |\Psi \cap \mathcal{A}_2|$. Then the symmetric difference $\mathcal{A}_1 \triangle \mathcal{A}_2$ is t-balanced.*

Proof. We partition $\mathcal{A}_1 \triangle \mathcal{A}_2$ as $(\mathcal{A}_1 \setminus \mathcal{A}_2) \uplus (\mathcal{A}_2 \setminus \mathcal{A}_1)$ and note that for any t^+-scope Ψ,

$$|\Psi \cap (\mathcal{A}_1 \setminus \mathcal{A}_2)| = |\Psi \cap \mathcal{A}_1| - |\Psi \cap (\mathcal{A}_1 \cap \mathcal{A}_2)| = |\Psi \cap \mathcal{A}_2| - |\Psi \cap (\mathcal{A}_1 \cap \mathcal{A}_2)| = |\Psi \cap (\mathcal{A}_2 \setminus \mathcal{A}_1)|.$$

□

The union of disjoint balanced families is, quite clearly, balanced as well. The next proposition shows that the union of non-disjoint weakly balanced families can be made weakly balanced by discarding their intersection.

Proposition 4.4. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two weakly t-balanced families. Then their symmetric difference $\mathcal{A}_1 \triangle \mathcal{A}_2 = (\mathcal{A}_1 \cup \mathcal{A}_2) \setminus (\mathcal{A}_1 \cap \mathcal{A}_2)$ is weakly t-balanced.*

Proof. Let Ψ be a t^+-scope. Then

$$|\Psi \cap (\mathcal{A}_1 \triangle \mathcal{A}_2)| = |(\Psi \cap \mathcal{A}_1) \triangle (\Psi \cap \mathcal{A}_2)| = |\Psi \cap \mathcal{A}_1| + |\Psi \cap \mathcal{A}_2| - 2|\Psi \cap \mathcal{A}_1 \cap \mathcal{A}_2|,$$

which is even. □

4.1 Structure of scopes

In this subsection we give some results relating scopes and their sub-scopes (i.e., scopes contained in them), and discuss applications to balanced families. Note that since any scope Ψ contains only comparable graphs, any sub-scope of Ψ can be describe simply as a set of the form

$$\{\, G \in \Psi \mid F_0 \subseteq G \,\}$$

for a graph F_0. We begin with a simple result showing that any scope can be uniformly covered by t-scopes, for *any* t.

Lemma 4.5. *For all positive integers $T \geq t$, any T-scope can be covered by t-scopes such that every graph is covered exactly $\binom{T}{t}$ times.*

Proof. Let Ψ be a T-scope with fixed part G_0, and consider the set of "partial completions" of G_0,

$$S := \{\, G_0 \uplus H_0 \mid \exists G_0 \uplus H \in \Psi \,.\, H_0 \subseteq H \text{ and } |H_0| = T - t \,\}.$$

Notice that for every member of S, the set of its supergraphs in Ψ is clearly a t-scope contained in Ψ. Furthermore, any graph $G_0 \uplus H \in \Psi$ is a supergraph of precisely those members $G_0 \uplus H_0$ of S for which $H_0 \subseteq H$. It follows that the aforementioned t-scopes (i.e., one for each member of S) cover every graph in Ψ exactly $\binom{T}{T-t} = \binom{T}{t}$ times. □

The next result shows that, in the definition of a t-balanced family, it suffices to consider t-scopes instead of t^+-scopes. Thus, if two families of comparable graphs are such that every graph of a certain size has the same number of supergraphs in each family, then the same is true for all graph of smaller size.

Proposition 4.6. *Let $\mathcal{A} = \mathcal{A}^+ \uplus \mathcal{A}^-$ be a family of graphs and suppose that for every t-scope Ψ it holds that $|\Psi \cap \mathcal{A}^+| = |\Psi \cap \mathcal{A}^-|$. Then $\mathcal{A}$ is t-balanced.*

Proof. Let Ψ be a t'-scope for some $t' \geq t$. By Lemma 4.5, Ψ can be covered by some t-scopes $\{\Psi_i\}_i$ such that each graph in Ψ is covered precisely $k := \binom{t'}{t}$ times. It follows that

$$k|\Psi \cap \mathcal{A}^+| = \sum_i |\Psi_i \cap \mathcal{A}^+| = \sum_i |\Psi_i \cap \mathcal{A}^-| = k|\Psi \cap \mathcal{A}^-|,$$

which completes the proof. $\square$

We record the relaxed (yet equivalent) definition of a balanced family obtained by combining Proposition 4.3 and Proposition 4.6.

Proposition 4.7. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two families of graphs. Suppose that for every t-scope Ψ it holds that $|\Psi \cap \mathcal{A}_1| = |\Psi \cap \mathcal{A}_2|$. Then the symmetric difference $\mathcal{A}_1 \triangle \mathcal{A}_2$ is t-balanced.*

Next we give a structure theorem for scopes, which would in particular imply that a result similar to Proposition 4.6 also holds for weakly balanced families (the proof of Proposition 4.6 fails to apply if, for some scope, the number of times a graph is covered is even). We will first need the following lemma showing that, among comparable graphs, it is always possible to choose subgraphs of bounded cardinality in a "canonical" way.

Lemma 4.8. *For all $t \geq 2$, every (partite) graph G of cardinality at least t contains a subgraph $\varphi(G) = \varphi_t(G)$, obtained by removing between t and $2t-2$ edges, such that for every graph G' that is comparable to G and contains $\varphi(G)$, it holds that $\varphi(G') = \varphi(G)$.*

Proof. We begin by defining, for every graph G of cardinality at least t, a subgraph $\mathrm{I}(G) = \mathrm{I}_t(G) \subseteq G$ of cardinality between t and $2t-2$. This will be used to define $\varphi(G)$ later. Fix an ordering of the vertex classes, and fix an ordering of the vertices in each vertex class by non-increasing degree. Viewing edges of a partite graph as sequences in a straightforward manner, consider the lexicographically-first edge e_1 in G, and let $k = k(G)$ denote the length of the shortest prefix of e_1 that is common to at most t edges. More formally—if $e_1 = \{v_1, \dots, v_d\}$, where v_i is from the ith vertex class, then $k(G)$ is the smallest j such that the *degree* of the prefix $\{v_1, \dots, v_j\}$, defined as $|\{\, e \in G \mid \{v_1, \dots, v_j\} \subseteq e \,\}|$, is at most t. Now, consider the lexicographically-first prefixes of length k in G such that the set of edges having either of those prefixes is of cardinality at least t (i.e., the degrees of the prefixes sum up to at least t). We define the subgraph $\mathrm{I}(G)$ as the latter set of edges, and claim that its cardinality is at most $2t-2$. Indeed, either $\mathrm{I}(G)$ is determined by a single length-k prefix of degree t, and then $\mathrm{I}(G)$ is of cardinality exactly t, or else every length-k prefix is of degree at most $t-1$, which implies that $\mathrm{I}(G)$ is of cardinality at most $2t-2$, as desired.

We now show that if G' is a rewiring of $\mathrm{I}(G)$ in G, obtained by replacing $\mathrm{I}(G)$ with a comparable F (so $G \setminus \mathrm{I}(G) = G' \setminus F$), then $\mathrm{I}(G') = F$. This would complete the proof, since for any graph G' that is comparable to G and contains $\varphi(G) := G \setminus \mathrm{I}(G)$, we have that $\varphi(G') = G' \setminus \mathrm{I}(G') = G \setminus \mathrm{I}(G) = \varphi(G)$, as required. Observe that, by definition, if $k = k(G) > 1$ then all edges in $\mathrm{I}(G)$ share a common prefix of length $k-1$. In other words, $\mathrm{I}(G)$ has a single vertex from each of the first $k-1$ vertex class. It follows that the same holds for F, as it is comparable to $\mathrm{I}(G)$. It is then not hard to see that G' has exactly the same prefixes of length k as G, which implies that $\mathrm{I}(G') = F$, completing the proof. $\square$

We use the above lemma to prove the following decomposition result for scopes.

Proposition 4.9. *For all $t \geq 2$, any t^+-scope can be partitioned into s_i-scopes where s_i are always at least t and at most $2t-2$.*

Proof. Let Ψ be a t'-scope where $t' \geq t$, with fixed part G_0. To define a partition of Ψ we need to assign to each member of Ψ a (sub-) scope, such that all these scopes are pairwise disjoint. Equivalently, we need to assign to each member of Ψ a subgraph whose every supergraph in Ψ is also assigned it. We thus assign to each member $G_0 \uplus H \in \Psi$ the subgraph $G_0 \uplus \varphi(H)$, where $\varphi(H) = \varphi_t(H) \subseteq H$ is defined as in Lemma 4.8; notice that the corresponding scope is an s-scope for some $t \leq s \leq 2t-2$, as needed. It is therefore enough to prove the following assertion. For every H' comparable to H, if H' contains $\varphi(H)$ then $\varphi(H') = \varphi(H)$. Since the latter follows from Lemma 4.8, we are done. $\square$

While we do not know if it is in fact enough to consider subgraphs of cardinality exactly t when considering weakly t-balanced families (as is the case for t-balanced families), we *can* deduce, as an immediate corollary of Proposition 4.9, that it is enough to consider t'-scopes for t' which is at most $2t-2$.

Corollary 4.10. *Let $\mathcal{A}$ be a family of graphs and suppose that for every t'-scope Ψ, where $t \leq t' \leq 2t-2$, it holds that $|\Psi \cap \mathcal{A}|$ is even. Then $\mathcal{A}$ is weakly t-balanced.*

Before continuing, let us give a simple example showing that while scopes can in general be partitioned into sub-scopes in many different ways, it is not true that every scope can be partitioned into t-scopes for *any* t.

Example 4.11. *Let Ψ be the set of 2-regular n-bounded bipartite graphs where $n \geq 3$. Notice that Ψ is a scope with an empty fixed part (i.e., an equivalence class of the comparable relation). Denote $T = 2n$ the cardinality of the graphs in Ψ (i.e., Ψ is a T-scope). We put $t = T-1$, and claim that Ψ cannot be partitioned into t-scopes. Consider any two t-scopes in Ψ, and let G_0 and G_0' denote their fixed parts, so that $|G_0| = |G_0'| = 1$. Notice that the intersection of the two scopes is precisely the collection of supergraphs in Ψ of $G_0 \cup G_0'$. The latter graph is of cardinality at most 2, and it clearly can be completed to a cycle C_{2n} on all vertices, which is a graph in Ψ. It follows that every two t-scopes in Ψ intersect. Since there can be no single t-scope that covers all of Ψ (there is no edge that appears in all graphs in Ψ, since for any two vertices u, v there is a cycle C_{2n} in which u and v are not adjacent), it follows that no partition of Ψ into t-scopes is possible.*

We now prove that if t is a power of two (i.e., $t = 2^a$ for some positive integer a), it is in fact enough to consider only t-scopes in the definition of weakly t-balanced families.

Corollary 4.12. *Let t be a power of two, and let $\mathcal{A}$ be a family of graphs. Suppose that for every t-scope Ψ it holds that $|\Psi \cap \mathcal{A}|$ is even. Then $\mathcal{A}$ is weakly t-balanced.*

Proof. Let Ψ be a t'-scope for $t' \geq t$. From Lemma 4.5 it is easy to see that if $k := \binom{t'}{t}$ is odd then $|\Psi \cap \mathcal{A}|$ is even. It follows from Lucas' Theorem that k is odd if and only if every digit in the base-2 expansion of t' is at least the corresponding digit of t. Since t is a power of 2, it holds that k is odd if t' is smaller than $2t$. By Corollary 4.10, we may assume that t' is at most $2t - 2 < 2t$, which completes the proof. $\square$

4.2 Amplification of balanceable graphs

Let G be a t-balanceable n-bounded p-partite graph. It is not hard to see that adding a new vertex to all edges of G (and also $n-1$ degree-0 vertices in the new vertex class) results in an n-bounded $(p+1)$-partite graph that is t-balanceable as well. Notice that the new graph has the same cardinality as G, and so cannot be used to deduce a larger complexity lower bound. In this subsection we show that if there is an explicit construction of a t-balanceable n-bounded p-partite graph G, then there is also an explicit construction of a t-balanceable n-bounded kp-partite graph G' of a much larger cardinality (provided $k > 1$).

Let us mention that the implied lower bound on the complexity of G' can be (rather easily) shown directly from the definition. However, here we prove the stronger claim that G' is balanceable, which for example implies (the same) lower bound for many more graphs using Theorem 3.11.

We begin with some simple results relating p-partite graphs and kp-partite graphs. Let $p \geq 2$, k, n be positive integers. Fix a bijection φ mapping $[n^k]$ to $[n]^k$. We also use φ to denote the bijection from $[n^k]^p$ to $[n]^{kp}$ naturally induced by φ, as well as to denote the induced bijection from n^k-bounded p-partite graphs to n-bounded kp-partite graphs, obtained by replacing every vertex v from an n^k-sized vertex class with a sequence $\varphi(v)$ of k vertices from n-sized vertex classes (more formally, $\varphi(v)$ is a set of vertices, each from a different n-sized vertex class).

Lemma 4.13. *If G and G' are comparable then so are $\varphi(G)$ and $\varphi(G')$.*

Proof. Let v be a vertex of $\varphi(G)$, and let S_v be set of vertices u of G such that $v \in \varphi(u)$. Denoting $\deg_H(v)$ the degree of v in a graph H, it is not hard to see that

$$\deg_{\varphi(G)}(v) = \sum_{u \in S_v} \deg_G(u) = \sum_{u \in S_v} \deg_{G'}(u) = \deg_{\varphi(G')}(v),$$

which completes the proof. $\square$

As a consequence of the above, we next show that scopes containing kp-partite graphs are very much related to scopes containing p-partite graphs. In what follows, we use φ to also denote a bijection between families of graphs, i.e., $\varphi(\mathcal{F}) = \{\, \varphi(G) \mid G \in \mathcal{F} \,\}$.

Lemma 4.14. *For every t, any t-scope containing n-bounded kp-partite graphs can be partitioned into images of t-scopes under φ.*

Proof. Notice that every n-bounded kp-partite graph is of the form $\varphi(G)$ for some n^k-bounded p-partite graph G. Let Ψ be a t-scope containing n-bounded kp-partite graphs, whose fixed part is $\varphi(G_0)$. We define an equivalence relation on Ψ by considering $\varphi(G), \varphi(G') \in \Psi$ equivalent if

G and G' are comparable. We next show that each equivalence class of the above relation is the image under φ of a t-scope, which would complete the proof.

Let $\varphi(G)$ be a member of an equivalence class. We need to show that for every graph G' that is comparable to G and contains G_0, it holds that $\varphi(G')$ is in the same equivalence class as $\varphi(G)$. Notice that by the definition of our equivalence relation, since G and G' are comparable, it suffices to show $\varphi(G') \in \Psi$. By Lemma 4.13 we have that $\varphi(G')$ is comparable to $\varphi(G) \in \Psi$, and since G' contains G_0, or equivalently, $\varphi(G')$ contains $\varphi(G_0)$, we conclude that $\varphi(G') \in \Psi$, as desired. □

We next conclude that φ preserves balance.

Corollary 4.15. *Let G be an n^k-bounded p-partite graph. Then if G is (weakly) t-balanceable then $\varphi(G)$ is (weakly, respectively) t-balanceable.*

Proof. Let $\mathcal{A}$ be a (weakly) t-balanced family containing G. It easily follows from Lemma 4.14 that $\varphi(\mathcal{A})$ is also a (weakly, respectively) t-balanced family, which completes the proof. □

Now, let $G = (G_n)_{n\in\mathbb{N}}$ be a (sequence of) n-bounded p-partite graph. For an integer $d \geq 2$, we define the *d-amplification* of G, assuming $d = kp$, as the n-bounded d-partite graph $\varphi(G_{n^k})$. If d is not a multiple of p, that is, $d = d' + c$ where d' is a multiple of p and $0 < c < p$, we define the d-amplification of G by adding c new vertices to all edges of the d'-amplification of G (together with $n-1$ degree-0 vertices in each of the c new vertex class). We next summarize the parameters of the above amplification.

Corollary 4.16. *Let G be an n-bounded p-partite graph of cardinality $m = m(n)$, and suppose that G is (weakly) t-balanceable. It holds that the d-amplification of G, which is an n-bounded d-partite graph of cardinality $m(n^{\lfloor d/p \rfloor})$, is also (weakly, respectively) t-balanceable.*

It follows that while the implied complexity lower bound on G is $r(n) := \left\lceil \frac{m(n)}{t-1} \right\rceil$, its d-amplification is of complexity at least $r(n^{\lfloor d/p \rfloor})$. Let us also note that the above corollary holds even if we allow t to depend on n as well.

Remark. Notice that amplifying p-partite graphs can imply only so much with regards to complexity lower bounds: since any p-partite graph is of complexity at most n^{p-1}, and in particular $r(n) := \left\lceil \frac{m(n)}{t(n)-1} \right\rceil \leq n^{p-1}$, it follows that its d-amplification implies a lower bound of $r(n^{\lfloor d/p \rfloor}) \leq (n^{\lfloor d/p \rfloor})^{p-1}$, which is no more than $n^{d-\lceil d/p \rceil}$ ($\ll n^{d-1}$ for $p < d$).

4.3 An equivalent definition of balanced families

In this subsection we present an equivalent definition of a balanced family in terms of local operations on graphs.

Definition 4.17. A *rewiring involution* is a graph operation that, given a graph G and a subgraph $H \subseteq G$, replaces H with a comparable graph—thus yielding a graph of the form $G' = (G \setminus H) \cup H'$ for some H' comparable to H—and such that applying it on H' *in* G' (i.e., applying it on G' and $H' \subseteq G'$) yields back G.

We next prove that a family is weakly balanced if and only if it is closed under a rewiring involution that has no fixed points, i.e., always yields a different graph (in the notation above, $G \neq G'$, or equivalently, $H \neq H'$).

Proposition 4.18. *A family of graphs $\mathcal{A}$ is weakly t-balanced if and only if there is a rewiring involution such that applying it on any subgraph of cardinality at least t in a graph of $\mathcal{A}$, yields a different graph of $\mathcal{A}$.*

Proof. The proof easily follows by observing that a rewiring involution as in the statement is equivalent to an involution without fixed points on the intersection of any t^+-scope with $\mathcal{A}$, and using the obvious fact that a set is of even cardinality if and only if there exists an involution on it without fixed points. □

Proposition 4.18 implies a recursive definition of weakly balanceable graphs as follows. For a rewiring involution f operating on subgraphs of cardinality at least t, say that a graph is f-balanceable if any application of f yields a different f-balanceable graph. Then clearly, a graph is weakly t-balanceable if and only if it is f-balanceable for some f.

Remark. Let f be a rewiring involution. Say that two graphs are f-equivalent if it is possible to obtain one graph from the other by repeatedly applying f. Since f is its own inverse, this is clearly an equivalence relation. Thus, any family closed under f is a disjoint union of equivalence classes of this relation.

We can extend the above definition to balanced families as follows (recall that by Proposition 4.6 it is enough to consider t-scopes, instead of t^+-scopes, in the definition of a t-balanced family).

Fact 4.19. *A family of graphs $\mathcal{A} = \mathcal{A}^+ \mathbin{\dot\cup} \mathcal{A}^-$ is t-balanced if and only if there is a rewiring involution such that applying it on any t-subgraph of a graph in $\mathcal{A}^+$ yields a graph in $\mathcal{A}^-$, and vice versa.*

It is in fact possible to define balanced families without an explicit mention of a bipartition of its members, as follows. Consider an undirected graph (i.e., not a hypergraph) whose vertices are the members of a weakly t-balanced family $\mathcal{A}$. Suppose that $\mathcal{A}$ is closed under the fixed-point-free rewiring involution f. For any two members G_1 and G_2 of $\mathcal{A}$, if applying f on a t-subgraph of G_1 gives G_2 then we add an (undirected) edge between the corresponding vertices in our graph. We allow multiple edges.

Notice that this graph has no loops, and is regular of degree $\binom{m}{t}$ if m is the cardinality of all the graphs in $\mathcal{A}$. We conclude the following.

Fact 4.20. *The above graph has no odd-length cycles (i.e., it is a bipartite graph) if and only if $\mathcal{A}$ is t-balanced.*

5 Constructing balanceable graphs

In this section we present some explicit examples of balanced families of graphs. Let us first establish a useful terminology. A *graph transposition* (or simply *transposition*) is the graph isomorphism[10] that transposes a pair of vertices (from the same vertex set). That is, a transposition of the two vertices v_1 and v_2 maps any graph G to the graph obtained by replacing every occurrence of v_1 in the edges of G with v_2, and vice versa (i.e., it swaps their labels).

Given a subgraph $H \subseteq G$, consider a transposition T on H that yields a graph disjoint from $G \setminus H$. In this case, T can naturally be viewed as an operation on G that leaves $G \setminus H$ unaltered, yielding the graph $G' = (G \setminus H) \cup T(H)$. Note that if $T(H)$ is not disjoint from $G \setminus H$, this operation would not result in a (simple) graph, as multiple edges would arise. Also note that G' is in general not isomorphic to G, even though $T(H)$ is isomorphic to H. The following intuitively-phrased observation (which is based on the notion of a rewiring involution from Definition 4.17) will be of use to us.

Fact 5.1. *A transposition of a pair of vertices in a subgraph H is a rewiring involution, provided the vertices are of the same degree in H, and that their choice is independent of H (their choice may depend on the rest of the graph).*

Combining the above fact with Proposition 4.18 immediately implies a criterion for showing that a given family is weakly balanced, which we dub the *Transposition Criterion* and use throughout this section.

Motivating example: irreducible graphs. The first example we present deals with graphs satisfying a property known as *irreducibility*. This property has been considered in the past, e.g., in the work of Mubayi [19] on so-called Turán densities. We show that irreducible graphs of cardinality m form a t-balanced family for t that is somewhat smaller than m (recall Example 3.8, in which m-balanced families are characterized). We remark that since t will be fairly close to m (i.e., $\left\lceil \frac{m}{t-1} \right\rceil$ is small), this example would not imply significant lower bounds.

A graph G is *irreducible* if every vertex v (appearing with positive degree) has a unique set of partial edges $N_G(v) = \{\, e \setminus \{v\} \mid v \in e \in G \,\}$. That is, for every two vertices u, v appearing in the same vertex class, $N_G(v) \neq N_G(u)$. We now show that irreducibility is a balanced property of graphs.

Claim 5.2. *Let G be an irreducible, regular n-bounded graph of cardinality m. Then G is weakly $(m-n+2)$-balanceable.*

Proof. Notice that we may assume $n \geq 2$ as otherwise the result is trivial. We show that the family of graphs isomorphic to G is weakly $(m-n+2)$-balanced. Let H be a subgraph of cardinality at least $m-n+2$ in a graph G' that is isomorphic to G. Since $G' \setminus H$ is of cardinality at most $n-2$, it must have at least two degree-0 vertices from each vertex class. Let u, v be two such vertices taken from an arbitrarily-chosen vertex class. Consider the graph obtained from G by transposing u and v in H. Since every edge that includes either u or v appears in H, this transposition is in fact a transposition on G'. This yields a graph isomorphic to G', and thus isomorphic to G as well. Furthermore, as G' is irreducible, it is clear that such a transposition

[10] A *graph isomorphism* maps any (partite hyper-) graph G to the graph obtained by permuting the labels of the vertices in each vertex class; more formally, it maps G to the graph $\{\, \{\pi(v_1), \ldots, \pi(v_d)\} \mid \{v_1, \ldots, v_d\} \in G \,\}$ where π is a permutation that maps each vertex to a vertex from the same vertex class.

results in a graph different from G'. Since u and v are of the same degree in H, the proof follows from the Transposition Criterion above. $\square$

Now that we have shown that any irreducible regular graph is weakly balanceable, we proceed to show it is in fact balanceable by specifying an appropriate bipartition of a weakly balanced family containing it.

Claim 5.3. *In the notation of Claim 5.2, G is $(m-n+2)$-balanceable.*

Proof. Let $V = \{v_1, \ldots, v_n\}$ be a vertex class of G. Since G is irreducible, every vertex $v_i \in V$ has a unique set of partial edges $N_G(v_i) = \{\, e \setminus \{v_i\} \mid v_i \in e \in G \,\}$. Let $\mathcal{A}$ be the family of graphs isomorphic to G that are obtained by permuting the vertices in V. Associate with every graph $G' \in \mathcal{A}$ the permutation on $[n]$ mapping i to j if and only if $N_{G'}(v_i) = N_G(v_j)$. Consider the bipartition of $\mathcal{A}$ according to the parity of the permutation associated with each member. Notice that transposing a pair of vertices of V in a graph from $\mathcal{A}$ always changes the parity of the associated permutation. Similarly to the proof of Claim 5.2, the rewiring involution that transposes (only) pairs of vertices from V shows that $\mathcal{A}$ is weakly $(m-n+2)$-balanced, and using Fact 4.19, we conclude that $\mathcal{A}$ is $(m-n+2)$-balanced, which completes the proof. $\square$

Notice that Claim 5.3 can alternatively be proved using Fact 4.20 by observing that a cycle in the graph mentioned there corresponds to a sequence of transposition permutations whose composition is the identity permutation; since the latter is an even permutation, it follows there are no odd cycles.

5.1 2-balanced families

The rest of the examples we present in this section are (weakly) t-balanced for the smallest possible t, namely $t = 2$. Such families have an especially interesting complexity implication, as follows. It is straightforward from definition that the complexity of any graph is at most its number of edges. Thus, the complexity of any (weakly) 2-balanceable graph is the maximum possible, that is, as high as its cardinality. In particular, showing that a graph is (weakly) 2-balanceable allows us to exactly determine its complexity.

As a special case of Claim 5.3, we next record that the family of matchings (i.e., graphs whose edges are pairwise disjoint) is 2-balanced.[11]

Claim 5.4. *Any matching is 2-balanceable.*

Proof. Notice that a graph is a matching if and only if every vertex of positive degree is of degree exactly 1; thus, the cardinality of a matching is equal to the number of vertices of positive degree from each vertex class. Furthermore, it is easy to see that any matching is an irreducible graph. The proof then immediately follows from Claim 5.3 $\square$

[11] Notice that when considering (weakly) t-balanced families, we may ignore all graphs of cardinality smaller than t, as the definition depends only on t^+-scopes, which do not contain such graphs. In particular, we may always discard any single-edge graph in a balanced family.

The largest n-bounded matchings are of course those without isolated vertices, that is, perfect matchings, and are of cardinality n. We conclude (the easy fact) that any n-bounded perfect matching is of complexity exactly n, over any field.

Remark. We can use the simple example above to deduce a (well-known) complexity lower bound much larger than n for d-partite graphs, assuming $d \geq 4$. By Corollary 4.16, the d-amplification of an n-bounded bipartite perfect matching, which is an n-bounded d-partite graph of cardinality $n^{\lfloor d/2 \rfloor}$, is of complexity exactly $n^{\lfloor d/2 \rfloor}$ over any field.

For bipartite graphs we can show the following: *Any* 2-balanced family of bipartite graphs is in fact a family of matchings (this discussion applies to weak balance as well). By Claim 3.3, the maximum degree of a t-balanceable bipartite graph is at most $t-1$ (as it may not contain a copy of $K(1,t)$; see also the discussion in the beginning of Section 4). In particular, any 2-balanceable bipartite graph must be of maximum degree 1, that is, a matching. Put differently—for bipartite graphs and $t = 2$, excluding the trivial obstruction of complete subgraphs already yields a (2-) balanced family. Notice that, since the complexity of any n-bounded d-partite graph is at most n^{d-1}, the lower bound of n we have obtained is optimal for bipartite graphs (thus, considering t-balanced families of bipartite graphs for $t > 2$ cannot yield better lower bounds).

Since the complexity of a graph is at least the complexity of any of its induced subgraphs, we may conclude the following.

Corollary 5.5. *Any graph that contains an induced matching of k edges is of complexity at least k over any field.*

In fact, using Theorem 3.11 we may deduce a (well-known) complexity lower bound of k over any field even if the size-k induced subgraph is not a matching by itself, but merely contains an odd number of perfect matchings.

5.2 Forests are balanceable

In this subsection we show that considering t-balanceable graphs even for $t = 2$ already allows us to obtain non-trivial lower bounds—that is, larger than n for n-bounded graphs—even when the graphs are 3-partite (as mentioned above, the complexity of n-bounded d-partite graphs for $d < 3$ is at most n). Specifically, we next give an explicit description of a 2-balanced family, and deduce lower bounds of up to $3n/2$ on the complexity of n-bounded 3-partite graphs (and more generally, d-partite graphs where $d \geq 3$). Let us mention that, unlike the previous examples, the balanced family we describe here contains, by definition, non-isomorphic graphs.

Interchanges. First, let us introduce terminology for the building block of rewiring operations. We call a graph transposition that is applied on a subgraph of cardinality 2 (i.e., a pair of edges) an *interchange*; the transposed vertices are said to be *interchanged*.

Example 5.6. *If we identify edges of d-partite graphs with d-tuples then, for the edges* $(3,1,4)$ *and* $(2,7,1)$*, interchanging the vertices* 3 *and* 2 *yields the edges* $(2,1,4)$ *and* $(3,7,1)$.

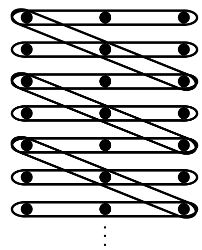

Figure 1: A 3-partite tree.

Remark. By a theorem of Ryser, any graph comparable to a graph G can be obtained by applying a (finite) sequence of interchanges on G. (A formal proof for bipartite graphs, phrased in a somewhat different language, was given in [24], where the term "interchange" first appeared in this context; we also mention that for non-bipartite graphs, the analogous notion is called a 2-*switch* and a similar result was shown, e.g., by Havel [12].)

Forests and trees. Before defining hypergraph forests, we first introduce a straightforward reduction. Given a 3-partite graph G, it will be useful to also view it as a graph $\bar{G}$ in the usual graph-theoretic sense (i.e., $\bar{G}$ is not a hypergraph). We construct $\bar{G}$, which is defined over the vertex set of G, as follows. For each edge $\{u, v, w\} \in G$, we arbitrarily choose one of its vertices, say v, and add to $\bar{G}$ the two incident edges $\{v, u\}$ and $\{v, w\}$. Notice that if G contains a pair of edges that intersect at two vertices then multiple edges may arise in $\bar{G}$ (we view a multiple edge as forming a length-2 cycle). Observe that consecutive vertices on a path in $\bar{G}$ are also adjacent (i.e., contained in the same edge) in G. It is worth mentioning that since G is a partite graph, consecutive vertices on a path in $\bar{G}$ belong to different vertex classes (in graph-theoretical terms, $\bar{G}$ is 3-colorable).

Finally, say that G is a *forest* (*tree*) if $\bar{G}$ is a forest (tree, respectively). Note that this is well defined since the existence of a cycle in $\bar{G}$ (and whether it is connected) is independent of the way we "break down" each edge of G.

Example 5.7. *The* 3*-partite graph illustrated in Figure 1 is an* n*-bounded tree, for any odd* n*.*

We next show that the family of 3-partite forests is weakly 2-balanced. By the Transposition Criterion, and since we may restrict to subgraphs of precisely two edges when considering weakly 2-balanced families (see Corollary 4.10), it is enough to show that every two edges of a 3-partite forest can be interchanged to give a different forest—where the choice of the pair of vertices to interchange depends only on the subgraph without the two edges.

Theorem 5.8. *Any* 3*-partite forest is weakly* 2*-balanceable.*

For the proof we will need to define a certain operation (in fact, rewiring involution) on forests which we call the *tree-swap* operation. Let e_1, e_2 be two edges in a 3-partite forest G. A *direct path* between e_1 and e_2 in $\bar{G}$ is a path whose endpoints s, t satisfy $s \in e_1$, $t \in e_2$, and are the only vertices on the path from these two edges. Observe that, since there is at most one

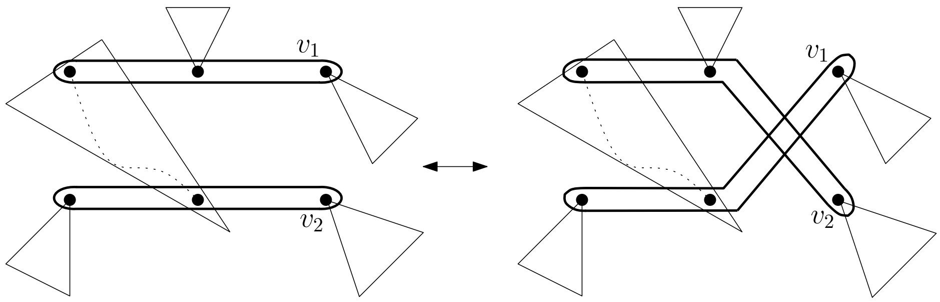

Figure 2: A pair of edges in 3-partite tree, before and after applying the tree-swap on them.

path between any two edges in a forest, there must exist two vertices from the same vertex class $v_1 \in e_1$ and $v_2 \in e_2$ neither of which on a direct path between e_1 and e_2. (Notice $v_1 \neq v_2$, as otherwise it forms a single-vertex direct path between the two edges.) We define the tree-swap operation so as to interchange v_1 and v_2 (which is well defined, being independent of the choices made in constructing $\bar{G}$); if there is more than one such pair of vertices, we choose arbitrarily. Note that the choice of v_1 and v_2 depends only on the degree list of G and on the subgraph of the forest without e_1 and e_2. Also, note that we are indeed allowed to apply the above interchange (i.e., by doing so we do not create multiple edges in G), since if one of the interchanged edges, say $(e_2 \setminus \{v_2\}) \uplus \{v_1\}$, already appears in the graph then, since this edge intersects e_2 at two vertices, $\bar{G}$ has a cycle, which is a contradiction.

Intuitively, the tree-swap operation can be thought to swap the subtrees "rooted" at each of the two interchanged vertices; see Figure 2 for an illustration.

Proof. We show that applying the tree-swap operation on any pair of edges, e_1 and e_2, in a 3-partite forest G yields a different forest G'. By the Transposition Criterion this would prove the result. Suppose that the tree-swap interchanges the vertices $v_1 \in e_1$ and $v_2 \in e_2$ (recall $v_1 \neq v_2$), and denote $e_1' = (e_1 \setminus \{v_1\}) \uplus \{v_2\}, e_2' = (e_2 \setminus \{v_2\}) \uplus \{v_1\} \in G'$ the interchanges edges.

First we claim $G' \neq G$. Indeed, equality would mean that G includes both e_2 and e_2' (notice $e_2 \neq e_2'$), which intersect at two vertices and thus imply a cycle in $\bar{G}$. Next, we show that $\bar{G}$ and $\bar{G}'$ have the same number of connected components. Since they obviously have the same number of vertices and edges, it would follow from a well-known characterization of forests that $\bar{G}'$ is a forest as well.

Observe that $\bar{G}$ and $\bar{G}'$ differ only in the edges of the form $\{v_1, x\}$ or $\{v_2, x\}$ where $x \in e_1 \cup e_2$. It follows that the paths contained in $\bar{G}$ and in $\bar{G}'$ are the same, except for those that include at least one of the above "differentiating" edges. Let v be a vertex on one of these paths in $\bar{G}$, or equivalently, in the connected component of either v_1 or v_2. Clearly, all other connected components are identical in the two graphs. Suppose that v is "directly" connected to a vertex $u \in e_1 \cup e_2$, in the sense that the (unique) path between v and u in $\bar{G}$ does not go through any other vertex from $e_1 \cup e_2$. Assuming $v \neq v_1, v_2$, this path does not include a "differentiating" edge, and so v and u remain connected in $\bar{G}'$. It follows that if $u = v_1$ (we can think of v in this case as belonging to the subtree "rooted" at v_1), or if $u \in e_2 \setminus \{v_2\}$, then it is the case that v is connected to v_1 in $\bar{G}'$. Otherwise, v is connected to v_2 in $\bar{G}'$. We conclude that the union of the connected components of v_1 and v_2 consists of the same vertices in $\bar{G}$ as in $\bar{G'}$; moreover, it consists of the same number of connected components—either one or two—in the two graphs, which completes the proof. $\square$

We remark that one can naturally extend the definition of a forest to d-partite graphs for $d > 3$, and similarly conclude that such a graph is weakly 2-balanceable. Another implication of the proof above is that the family of 3-partite forests with any given number of connected components is weakly 2-balanced as well. In particular, the family of 3-partite trees is weakly 2-balanced.

We may conclude that for a graph to have complexity any lower than the trivial upper bound—its cardinality—over a field of characteristic 2 (or the integers), it must contain at least one cycle. In particular, any graph containing an induced forest of cardinality k is of complexity at least k over any field of characteristic 2. By Theorem 3.11, one can use the fact that the family of forests is weakly 2-balanced to deduce the following: For a graph to be of complexity *smaller* than k over a field of characteristic 2, it must contain an even number of size-k forests of each possible degree list.

Remark. It follows from the fact that the family of forests (or trees) is weakly 2-balanced that the number of non-trivial (i.e., containing more than one edge) 3-partite forests (or trees) is even, regardless of the underlying vertex set. In fact, it follows that the number of such forests (or trees) with any given degree list is even as well.

We in particular obtain the following corollary concerning trees, which are the largest forests.

Corollary 5.9. *Any n-bounded 3-partite tree is of complexity exactly $(3n-1)/2$ over any field of characteristic 2.*

Proof. Let G be a 3-partite tree on N vertices. Clearly, the number of edges in $\bar{G}$ is exactly $N-1$. Since $\bar{G}$ has two edges per every edge of G, we conclude that the cardinality of G is exactly $(N-1)/2$. In particular, if G has $N = 3n$ vertices then it is of cardinality $(3n-1)/2$. Since any 3-partite tree is weakly 2-balanceable, the result immediately follows from Theorem 3.10. $\square$

It is not hard to see that n-bounded 3-partite trees exist; see, e.g., Example 5.7.

Say that a 3-partite tree is *binary* if its vertices are of degree at most 2 (i.e., each vertex is contained in at most two edges). We next show that the family of binary 3-partite trees, which is a subfamily of the above-mentioned weakly balanced family of forests, is in fact balanced.

Theorem 5.10. *Any binary 3-partite tree is 2-balanceable.*

Proof. We prove the result by giving an analogue for the parity of a permutation (which was used to define the bipartition of the family of matchings via Claim 5.3). The proof shall follow these steps: first, we define for every binary 3-partite tree G an ordering $\lessdot$ on the vertices $v_1, \ldots, v_N$, and denote π_G the permutation on $[N]$ mapping i to j if and only if v_i is in the jth position according to $\lessdot$. We then bipartition the family of binary 3-partite trees according to the parity of the associated permutation, from which the proof would follow, using Fact 4.19, by showing that this parity changes whenever we apply the tree-swap operation.

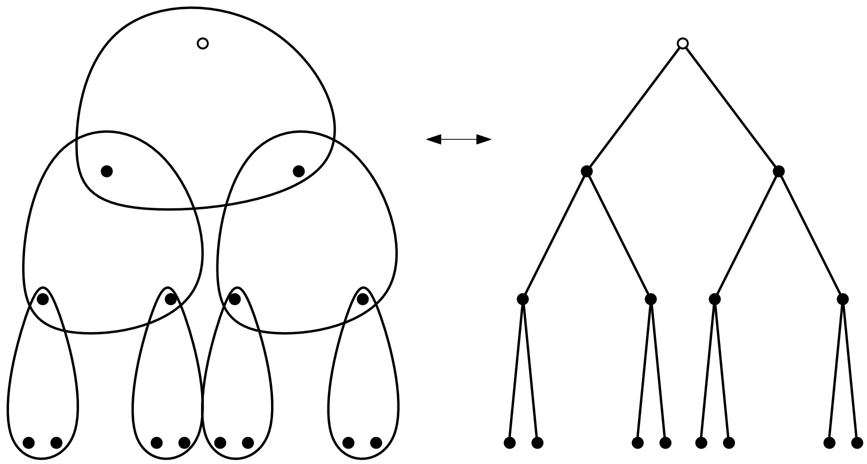

Figure 3: A rooted binary 3-partite tree (left) and the corresponding rooted binary tree (right).

Given a binary 3-partite tree G, we next show that $\bar{G}$ can be made into a binary tree T_G, in which every non-leaf vertex has exactly two children. Fix an arbitrary vertex of degree 1, referred to as the root (henceforth, all trees are implicitly rooted). To construct T_G, we first "break down" the unique edge of G containing the root in such a way that both new edges are adjacent to the root[12]; the remaining two vertices are referred to as children (of the root). We obtain T_G by continuing in the same manner for (the subtrees rooted at) each of the two children.[13] Moreover, we consistently distinguish between left and right children in T_G as follows. For a parent vertex from the vertex class numbered $i \in \{0, 1, 2\}$, its child from the vertex class numbered $i - 1 \bmod 3$ is designated as its left child (and the other child is of course designated as the right child). We mention that this construction may in fact be reversed so as to obtain G from T_G (up to a circular relabeling of the vertex classes); this bijection is illustrated in Figure 3.

Let us define an ordering $\lessdot$ on the vertices of a binary 3-partite tree G, which is what is known as the *inorder* of the tree T_G. In more detail, let u, v be two vertices of G; we denote $u \lessdot v$, and say that u is *to the left* of v, if in the tree rooted at the common ancestor of u and v in T_G, either u is in the left subtree, or (in case u itself is the common ancestor) v is in the right subtree. See Figure 4 for an illustration.

Let G be a binary 3-partite tree. To prove the result, we need to show that when applying the tree-swap operation on any pair of edges in G, the parity of the permutation π_G changes, or equivalently, the number of *inversions*—pairs of vertices whose inorder changes—is odd. Suppose that applying the tree-swap operation on $\{e_1, e_2\} \subseteq G$ interchanges the vertices $v_1 \in e_1$ and $v_2 \in e_2$. We consider two possible cases in T_G.

Case 1: v_1 is not an ancestor of v_2 and vice versa. We claim that both v_1 and v_2 are children in their respective edges. Indeed, if v_1, say, were the parent in e_1 then one of the children in

[12]Formally, if v is the root and $\{u, v, w\} \in G$, we add to T_G the two edges $\{u, v\}$ and $\{v, w\}$.

[13]Formally, denoting T_H^x the binary tree obtained from the 3-partite H having vertex x as its root, and denoting $G' = G \setminus \{u, v, w\}$, we have $T_G = T_G^v = \{\{u, v\}, \{v, w\}\} \cup T_{G'}^u \cup T_{G'}^w$, where T_H^x is empty if x is of degree 0.

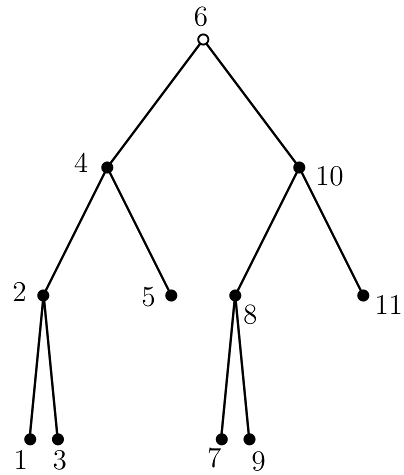

Figure 4: The inorder of a tree.

e_1 would have to be an endpoint of the direct path between e_1 and e_2, implying that v_1 is an ancestor of v_2, in contradiction to our assumption.

Intuitively, one can think of the tree-swap operation in this case as swapping the subtrees in T_G rooted at each of the interchanged vertices. Let us be more formal. Consider the set of vertices below a given vertex v, that is, the vertices in the subtree rooted at v, with v removed. Note that the number of vertices in any (sub-) tree is odd, and so there is an even number of vertices below any vertex. One way to prove that the parity of π_G changes after applying the tree-swap is by observing that the tree-swap acts on π_G by "flipping" two non-overlapping intervals of odd length—corresponding to the sets of vertices below v_1 and below v_2—which can be seen to be a permutation of odd parity. We next give an alternative proof, which will be useful when considering Case 2.

Denote T^* the set of vertices below v_1. The inorder between any two vertices in T^* does not change after applying the tree-swap, as no edge in their paths to their common ancestor is altered, which means that any inversion involving a vertex from T^* must involve exactly one such vertex. Now, let $u \notin \{v_1, v_2\}$ be a vertex outside of T. Note that the common ancestor of u and any vertex from T^* is exactly the common ancestor of u and v_1, while after applying the tree-swap, it is exactly the common ancestor of u and v_2. It follows that the number of inversions involving such vertices u is a multiple of $|T^*|$, and thus even. Furthermore, assuming $v_1 \lessdot v_2$, it is easy to see that a vertex $x \in T^*$ changes its inorder with respect to v_1 if and only if x is in the right subtree of v_1, which is further equivalent to x changing its inorder with respect to v_2. Hence the number of inversions involving $u \in \{v_1, v_2\}$ and a vertex from T^* is again even. We conclude that when considering the parity of the number of inversions, there is no loss of generality in removing the vertices below v_1 from T_G. Applying the same argument for the set of vertices below v_2, we thus may assume v_1 and v_2 to be leaves in T_G. The tree-swap then reduces to a transposition of v_1 and v_2 (which we may think of as simply switching the labels of v_1 and v_2). Applying a graph transposition has the effect of applying a transposition permutation on π_G, thus changing its parity, as desired.

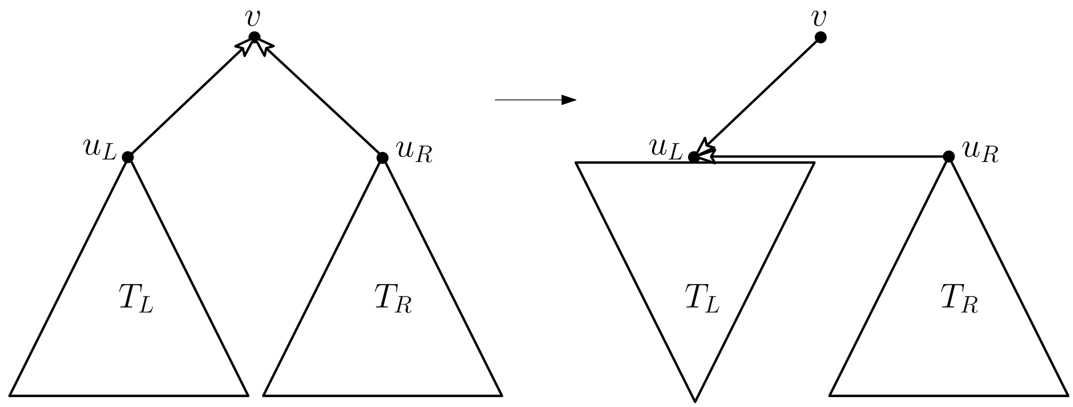

Figure 5: T_G with edges pointing toward the root, before and after re-rooting.

Case 2: v_1 is an ancestor of v_2. We claim that v_1 is the parent in e_1. Indeed, if v_1 were a child then, since it is an ancestor of v_2, it would have been an endpoint of the direct path between e_1 and e_2, which contradicts the definition of a tree-swap. Also, by the same argument, v_2 is a child in e_2. Notice that v_1 remains a parent after applying the tree-swap operation, as its path to the root is unaltered.

No inversion can involve a vertex outside the subtree T rooted at v_1, as no edge in a path from such a vertex to an ancestor is altered. Thus, we may consider only pairs of vertices that are both from T. In other words, if G' denotes the 3-partite graph corresponding to T, we may consider only inversions due to applying the tree-swap on G'. As in the previous case, we may assume that v_2 is a leaf. Since v_1 is of degree 1 in G' (albeit of degree 2 in G, unless it is the root of G), it is not hard to see that applying the tree-swap operation is equivalent to *transposing* v_1 and v_2 in G', followed by rooting G' at v_1 which is now a leaf. Similarly to the previous case, the first step results in an odd number of inversions. It therefore remains to show that the second step results in an even number of inversions, and this is proved in the following claim (which we remark is heavily dependent on the way we distinguish between left and right children).

Claim 5.11. *The number of inversions due to a re-rooting of a binary* 3*-partite tree at any leaf (i.e., taking any degree-*1 *vertex as the new root) is even.*

Proof. Let G be a binary 3-partite tree. Let v denote the root of G, and let u_L, u_R denote its left and right children in T_G, respectively. Assume that the future root is in the left subtree of v (the other case is treated similarly). Let T_L, T_R denote the set of vertices in the left and right subtrees of v in T_G, respectively. We next specify a partition of the set of pairs of vertices $\{x, y\}$ that involve at least one vertex outside of T_L, and show that the number of inversions in each part is even. This would complete the proof since by applying our claim by induction on the tree rooted at u_L (which includes the future root), it would follow that the number of inversions involving both vertices from T_L is also even. The partition is as follows.

- $x, y \in T_R$. There are no inversions involving both vertices from T_R, as the common ancestor of x and y is in T_R.

- $x \in T_R$, $y \in \{v, u_L\}$. After rooting, u_L becomes the parent of v and u_R. Because of the circular ordering of the vertex classes, it is not hard to see that v turns into a right child, and u_R into a left child (see Figure 5). Hence any vertex in T_R is now to the left of both v and u_L. Therefore, the number of inversions here is $2|T_R|$.
- $x = v, y = u_L$. By the previous argument, v becomes the right child of u_L after rooting, which means u_L remains to the left of v.
- $x \notin T_L$, $y \in T_L \setminus \{u_L\}$. A vertex in $T_L \setminus \{u_L\}$ changes its inorder with respect to a vertex in T_R if and only if it does so with respect to v, since, after rooting, the common ancestor of v and any vertex from T_R is u_L. It follows that the numbers of inversions here is a multiple of the number of vertices outside T_L, and thus even.

Note that the above is indeed a partition of the pairs involving a vertex outside T_L, which follows from the fact that the first three parts (items) cover exactly the pairs both of whose vertices are outside T_L, as well as the pairs that consist of u_L and a single vertex outside T_L. □

Since we have shown that any application of the tree-swap operation changes the parity of the number of inverted pairs, the proof is complete. □

Using Theorem 3.9 we can thus conclude that any binary 3-partite tree of cardinality m is of complexity m over *any* field. In particular, we may deduce the following corollary for n-bounded binary 3-partite trees (e.g., the tree in Example 5.7), whose proof is essentially the same as that of Corollary 5.9.

Corollary 5.12. *Any n-bounded binary 3-partite tree is of complexity exactly $(3n-1)/2$ over any field.*

A limitation of 2-balancedeness. We close this section with a proof showing that the above results are essentially tight for 2-balanceable 3-partite graphs, in the sense described below. Let G be a 3-partite graph. Recall that we may view $\bar{G}$ as being properly vertex-colored by 3 colors, corresponding to the partition into vertex classes. Also, observe that if $\bar{G}$ contains a cycle then, by possibly "breaking down" the edges of G differently, we may construct $\bar{G}$ so as to contain a cycle with at most one edge $\{u, v\}$ per every edge $\{u, v, w\}$ of G. We refer to such a cycle as a cycle of G (thus, intuitively, a cycle of G uses each edge of G at most once). We call a cycle of G *diverse* if it has a vertex with each color.

Proposition 5.13. *No 3-partite graph having a diverse cycle is (weakly) 2-balanceable.*

Proof. Let $\mathcal{A}$ be a weakly 2-balanced family of 3-partite graphs. We claim that if there is a member of $\mathcal{A}$ having a diverse cycle of length $k \geq 3$ then there must also be a member of $\mathcal{A}$ having either a diverse cycle of length $k-1$ or a cycle of length 2. Notice that a cycle of length 2 is a complete partite graph (it is synonymous with a 2-cluster $K(1, 1, 2)$), and so no member of a weakly 2-balanced family of 3-partite graphs has it. Therefore, proving the above claim would complete the proof.

Let $G \in \mathcal{A}$ have a diverse cycle C of length at least 4. Observe that there must be three consecutive vertices of C that are colored with a different color each. Indeed, going over the vertices in the cycle, starting from an arbitrary vertex, until the first time vertices of all three colors have been encountered clearly yields a vertex triple as required.

We claim that we can choose the above vertex triple so that the middle vertex is not uniquely colored, i.e., is not the only vertex of C with its color. Indeed, if there is a uniquely colored vertex v of C, then the rest of the vertices of C are properly colored using only two colors. Let u be a vertex adjacent to v in C. Note that u is not uniquely colored, since otherwise the rest of the vertices—of which there are at least two—would be properly colored using a single color. Thus, by taking the three consecutive vertices of C whose middle vertex is u we obtain a vertex triple as required.

Let $H \subseteq G$ denote the pair of edges in G corresponding to the aforementioned vertex triple. Since the two edges of H intersect, there is only a single other rewiring $G' \neq G$ of H in G (or else G already contains one of the "cross edges", and so has a length-2 cycle and we are done). Because $\mathcal{A}$ is (weakly) 2-balanced, it follows that G' is a member of $\mathcal{A}$ as well. It is not hard to see that G' has a cycle C' of length $k-1 \geq 3$, bypassing the vertex u in which the two edges of H intersect, which follows from the fact that the two vertices adjacent to u in C are from two different vertex classes. Furthermore, since C' has the same vertices as C except for u, and since u is not uniquely colored in C, it follows that C' is also diverse. Finally, by a similar argument, if G has a cycle of length 3, then $\mathcal{A}$ must contain a member having a cycle of length 2, and the proof is complete. $\square$

We next derive a simple upper bound on the cardinality of a 2-balanceable 3-partite graph.

Corollary 5.14. *Any (weakly) 2-balanceable n-bounded 3-partite graph is of cardinality at most $3n$.*

Proof. Let G be a weakly 2-balanceable 3-partite graph, and suppose it has a diverse cycle C. We claim that for any edge $e := \{u, v, w\} \in G$ corresponding to an edge $\{u, v\}$ of C, the vertex w must be of degree 1 in G. Indeed, if G has another edge e' containing w, then, similarly to the above proof, the subgraph $\{e, e'\} \subseteq G$ has a single other rewiring G' in G; it is not hard to see that G' has a diverse cycle, contradicting the fact that it is weakly 2-balanceable.

Consider all the cycles of G, and denote $H \subseteq G$ the corresponding edges of G. Clearly, the remaining edges of G form a forest. Observe that the forest has at most $3n - |H|$ vertices, which follows from the fact that every edge of H includes a vertex that appears only in H (i.e., of degree 1 in G). We conclude that the cardinality of G is at most $|H| + (3n - |H| - 1)/2 = (3n + |H| - 1)/2 \leq (6n - 1)/2$. $\square$

6 Proofs of the Main Theorems

In this section we prove the Main Theorems, Theorem 3.9 and Theorem 3.10, as well as Theorem 3.11. We start by naturally associating any family of graphs $\mathcal{A}$ with a polynomial $Q_{\mathcal{A}}$ that acts on tensors. We then use our notion of balance—we show that if $\mathcal{A}$ is a balanced family then $Q_{\mathcal{A}}$ vanishes on all low-rank tensors. The desired complexity lower bounds are obtained by showing that $Q_{\mathcal{A}}$ does not vanish on, e.g., the adjacency tensor of any member of $\mathcal{A}$.

Let us be more formal. Fix positive integers $d \geq 2$, n, and let $\mathbb{F}$ be an arbitrary field. Let $Q \in \mathbb{F}[X]$ be a polynomial in the variables $X = (X_e)_{e\in[n]^d}$; notice we may interpret Q as acting on d-tensors over $\mathbb{F}$. Given a family $\mathcal{A}$ of n-bounded d-partite graphs, where we naturally identify each such graph with a subset of $[n]^d$, consider the multilinear polynomial $Q_{\mathcal{A}} := \sum_{G\in\mathcal{A}} \prod_{e\in G} X_e$. Below we describe the behavior of multilinear polynomials on low-rank tensors, and then deduce that $Q_{\mathcal{A}}$ vanishes on all them if $\mathcal{A}$ is a balanced family. But first, we will need the following characterization of tensor rank.

Let r be a nonnegative integer. Define the mapping $\mathbb{L}^r \in \mathbb{F}^{rdn} \to \mathbb{F}^{n^d}$ by

$$\mathbb{L}^r : \alpha = (\alpha^i_{j,k})_{i\in[r],j\in[d],k\in[n]} \mapsto \Big(\sum_{i=1}^{r}\prod_{j=1}^{d}\alpha^i_{j,e_j}\Big)_{e=(e_1,\ldots,e_d)\in[n]^d}.$$

Lemma 6.1. *A d-tensor $A = (a_e)_{e\in[n]^d}$ over a field $\mathbb{F}$ is of rank at most r if and only if it lies in the image of $\mathbb{L}^r$.*

Proof. By definition, A is of rank at most r if there exist r tensor products of vectors in $\mathbb{F}^n$, $v^1_1 \otimes \cdots \otimes v^1_d, \ldots, v^r_1 \otimes \cdots \otimes v^r_d$, whose sum is A. Denoting $v^i_j = (\alpha^i_{j,1}, \ldots, \alpha^i_{j,n})$ where $\alpha^i_{j,k} \in \mathbb{F}$, the above is equivalent to the fact that for every $e = (e_1, \ldots, e_d) \in [n]^d$,

$$a_e = \alpha^1_{1,e_1}\cdots\alpha^1_{d,e_d} + \cdots + \alpha^r_{1,e_1}\cdots\alpha^r_{d,e_d},$$

which completes the proof. □

It follows from the above characterization that the restriction of $Q \in \mathbb{F}[X]$ to tensors of rank at most r is a polynomial as well, namely, $Q \circ \mathbb{L}^r := Q(\mathbb{L}^r(\alpha))$. We now show that this composed polynomial can in fact be described using the terminology of scopes.

Proposition 6.2. *Let $t > 1$ and m be positive integers, and let $\mathbb{L}^r(\alpha)$ be defined as above where $r < \left\lceil \frac{m}{t-1} \right\rceil$. Let $Q \in \mathbb{F}[X]$ be a multilinear polynomial in the variables $X = (X_e)_{e\in[n]^d}$,*

$$Q = \sum_G C_G \prod_{e\in G} X_e\,,$$

where the summation is over all graphs $G \subseteq [n]^d$ of cardinality at least m. Then for every monomial in α there are t^+-scopes $\{\Psi_i\}_i$ so that its coefficient in $Q \circ \mathbb{L}^r \in \mathbb{F}[\alpha]$ is

$$\sum_i \sum_{G\in\Psi_i} C_G\,.$$

Proof. Consider a monomial of Q,

$$\prod_{e\in G} X_e.$$

Evaluating the monomial on $\mathbb{L}^r(\alpha)$ gives

$$\prod_{e\in G} \mathbb{L}^r(\alpha)_e = \prod_{e\in G} \sum_{i=1}^{r} \prod_{j=1}^{d} \alpha^i_{j,e_j}.$$

Expanding the multiplication by distributivity, it is not hard to see that each resulting term corresponds to a coloring $\mathcal{G}$ of G with colors from $\{1,\dots,r\}$, or an r*-coloring*, and is of the form

$$p(\mathcal{G}) := \prod_{i=1}^{r} \prod_{e\in\mathcal{G}^i} \prod_{j=1}^{d} \alpha^i_{j,e_j},$$

where $\mathcal{G}^i$ denotes the color-i subgraph in $\mathcal{G}$. Denoting $\overline{\mathcal{G}}$ the underlying graph of the colored graph $\mathcal{G}$, we deduce that

$$Q(\mathbb{L}^r(\alpha)) = \sum_G C_G \prod_{e\in G} \mathbb{L}^r(\alpha)_e = \sum_G C_G \sum_{\mathcal{G}:\overline{\mathcal{G}}=G} p(\mathcal{G}),$$

where the innermost sum in the right hand side is over all r-colorings of G. Equivalently,

$$Q(\mathbb{L}^r(\alpha)) = \sum_{\mathcal{G}} C_{\overline{\mathcal{G}}} \cdot p(\mathcal{G}),$$

where the sum is over all r-colored graphs of cardinality at least m.

Say that two r-colored graphs are *comparable* if for every color $i\in[r]$, their color-i subgraphs are comparable.

Lemma 6.3. *The r-colored graphs $\mathcal{G}$ and $\mathcal{H}$ are comparable if and only if $p(\mathcal{G}) = p(\mathcal{H})$.*

Proof. By definition, $p(\mathcal{G}) = p(\mathcal{H})$ if and only if for every $1\le i\le r$, and for every $1\le j\le d$,

$$\prod_{e\in\mathcal{G}^i} \alpha^i_{j,e_j} = \prod_{e\in\mathcal{H}^i} \alpha^i_{j,e_j},$$

where $\mathcal{G}^i$, $\mathcal{H}^i$ denote the color-i subgraphs of G and H, respectively. By commutativity, equality holds in the above equation if and only if there is the equality of *multisets*

$$\{\, e_j \mid e\in\mathcal{G}^i \,\} = \{\, e_j \mid e\in\mathcal{H}^i \,\},$$

or equivalently, if and only if the same vertices, and with the same degrees, appear in the jth vertex class of $\mathcal{G}^i$ and $\mathcal{H}^i$. Since the latter is the same as saying that the graphs $\mathcal{G}^i$ and $\mathcal{H}^i$ are comparable, the proof is complete. □

Thus, each monomial of $Q(\mathbb{L}^r(\alpha))$ corresponds to an equivalence class Φ of the comparable equivalence relation on r-colored graphs of cardinality at least m, and its coefficient is

$$\sum_{\mathcal{G}\in\Phi} C_{\overline{\mathcal{G}}}\,.$$

Note that for every color i, the color-i subgraphs of the members of Φ are comparable graphs, and in particular, are of the same cardinality. This also implies that all members of Φ are of the same cardinality. Observe that in any coloring of a graph of cardinality at least m by $r < \left\lceil \frac{m}{t-1} \right\rceil$ colors there must be a monochromatic subgraph of cardinality at least t. Suppose there are $t' \geq t$ edges colored $c \in [r]$ in all members of Φ. Define an equivalence relation $\sim$ on Φ as follows. Say that two colored graphs satisfy $\sim$ if they are equal (i.e., same edges and identically colored) except, perhaps, for their c-colored edges. This clearly implies a partition of Φ into equivalence classes. It is not hard to see that for any such equivalence class, the underlying graphs of its members form a t'-scope.

We conclude that for each monomial of $Q(\mathbb{L}^r(\alpha))$ there are t^+-scopes $\{\Psi_i\}_i$, corresponding to the equivalence classes of $\sim$, such that its coefficient is

$$\sum_i \sum_{G \in \Psi_i} C_G,$$

as desired. $\square$

We deduce the following for the polynomial $Q_{\mathcal{A}}$ (recall $Q_{\mathcal{A}} = \sum_{G \in \mathcal{A}} \prod_{e \in G} X_e$) when $\mathcal{A}$ is balanced.

Corollary 6.4. *Let $\mathcal{A}$ be a weakly t-balanced family whose graphs are of cardinality at least m, and fix $r < \left\lceil \frac{m}{t-1} \right\rceil$. Then $Q_{\mathcal{A}} \circ \mathbb{L}^r \in \mathbb{F}[\alpha]$ is the zero polynomial if $\mathbb{F}$ is of characteristic 2.*

Proof. It follows from Proposition 6.2 that each monomial of $Q_{\mathcal{A}} \circ \mathbb{L}^r$ has coefficient of the form $\sum_{\Psi} |\Psi \cap \mathcal{A}|$, where the sum is over some t^+-scopes. Hence all these coefficients are even, and so vanish assuming characteristic 2. $\square$

When $\mathcal{A}$ is balanced, we may consider a variant of $Q_{\mathcal{A}}$ which would allow us to deduce a stronger result.

Corollary 6.5. *Let $\mathcal{A} = \mathcal{A}^+ \cup \mathcal{A}^-$ be a t-balanced family whose graphs are of cardinality at least m, and fix $r < \left\lceil \frac{m}{t-1} \right\rceil$. Denote $Q'_{\mathcal{A}} = Q_{\mathcal{A}^+} - Q_{\mathcal{A}^-}$. Then $Q'_{\mathcal{A}} \circ \mathbb{L}^r \in \mathbb{F}[\alpha]$ is the zero polynomial, for any field $\mathbb{F}$.*

Proof. By Proposition 6.2, the coefficients in $Q'_{\mathcal{A}} \circ \mathbb{L}^r$ are of the form $\sum_{\Psi} |\Psi \cap \mathcal{A}^+| - |\Psi \cap \mathcal{A}^-|$, where the sum is over some t^+-scopes. Hence all these coefficients are zero (for any field $\mathbb{F}$). $\square$

Intuitively, any nonzero function that vanishes on $\mathbb{L}^r$ carries a lot of information on tensors of rank at most r. In particular, if it is a multilinear, homogeneous polynomial then any monomial in its support corresponds to a tensor of rank larger than r, which proves the Main Theorems as follows.

Proof of Theorem 3.10. Let G be a graph that is contained in a weakly t-balanced family $\mathcal{A}$. We assume without loss of generality that the graphs in $\mathcal{A}$ are all of the same cardinality m (recall Fact 3.5). Evaluate $Q_{\mathcal{A}} = Q_{\mathcal{A}}(X)$ on (the adjacency tensor of) G, by setting $X_e = 1$ if $e \in G$, and $X_e = 0$ otherwise. It is easy to verify that, since $Q_{\mathcal{A}}$ is both homogeneous and multilinear, the above evaluates to 1. This completes the proof since $Q_{\mathcal{A}}$ vanishes, over any field of characteristic 2, on (the adjacency tensors of) all graphs of complexity smaller than $\left\lceil \frac{m}{t-1} \right\rceil$. $\square$

Note that evaluating the above-mentioned (homogeneous, multilinear) $Q_{\mathcal{A}}$ on an arbitrary graph G' yields the number of graphs from $\mathcal{A}$ that are subgraphs of G'; hence, if this number is odd then G' is also of complexity at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field of characteristic 2. This proves the first part of Theorem 3.11.

Proof of Theorem 3.9. Let G be a graph that is contained in a t-balanced family $\mathcal{A}$ whose graphs are, without loss of generality, all of the same cardinality m. Let $Q'_{\mathcal{A}}$ be the polynomial given in Corollary 6.5. Then, similarly to the previous proof, $Q'_{\mathcal{A}}$ evaluates on G to either 1 or -1, while $Q'_{\mathcal{A}}$ vanishes, over any field $\mathbb{F}$, on all graphs of complexity smaller than $\left\lceil \frac{m}{t-1} \right\rceil$. $\square$

Similarly to the previous comment, for a balanced family $\mathcal{A} = \mathcal{A}^+ \cup \mathcal{A}^-$, if a graph G' contains a different number of subgraphs from $\mathcal{A}^+$ and $\mathcal{A}^-$ then G' is of complexity at least $\left\lceil \frac{m}{t-1} \right\rceil$ over any field of characteristic zero—and also over any field of characteristic p, assuming the difference between these two numbers is not a multiple of p; in particular, if the difference is 1 then the conclusion holds over any field. This completes the proof of Theorem 3.11.

6.1 Extensions

The above proofs of the Main Theorems in fact imply somewhat stronger results, as follows.

- The same lower bounds hold also over any commutative rings (instead of fields) with multiplicative identity $1 \neq 0$.
- The same lower bounds hold even if we replace, in the definition of an adjacency tensor, any unity entry by any nonzero entry. More formally, suppose we associate any n-bounded d-partite graph G with some d-tensor $(a_e)_{e\in[n]^d}$ over a field $\mathbb{F}$ where a_e is nonzero if and only if $e \in G$ (identifying G with a subset of $[n]^d$). Then for any weakly t-balanceable graph of cardinality m, the associated tensor is of rank at least $\left\lceil \frac{m}{t-1} \right\rceil$ assuming $\mathbb{F}$ is of characteristic 2. Moreover, if G is balanceable then the same holds for any field $\mathbb{F}$.
- The same lower bounds hold also for the *border rank* over an algebraically closed field. The border rank of a tensor A is the smallest r so that every polynomial that evaluates to the zero polynomial on $\mathbb{L}^r$ also vanishes on A (this is the same as saying that the tensor lies in the so-called *Zariski closure* of the image of $\mathbb{L}^r$). Clearly, the border rank of a tensor is at most its rank, and so proving border rank lower bounds is harder. We mention it is known that the border rank of a tensor A can be interpreted as the smallest rank of tensor that "approximates" A; that is, showing that A is of border rank larger than r also implies that no sequence of rank-r tensors converges to A (see also [15], extending the work of Alder [2]).

Remark. Interpreting the balanced families in Section 5 algebraically, that is, as the support of a polynomial, it is not hard to see that the family of bipartite perfect matchings, shown to be

2-balanced in Claim 5.4, corresponds to the determinant polynomial. Furthermore, the family of trees, shown to be weakly 2-balanced in Theorem 5.8 (see also Theorem 5.10), and which consist of d-partite graphs for $d > 2$ is, when interpreted algebraically, reminiscent of the polynomial implied by the well-known Matrix-Tree Theorem.

Although we do not do so in this paper, it may also be interesting to consider the following alternative notions of vanishing of $Q \circ \mathbb{L}^r$ for a polynomial Q.

- Over a finite field of odd characteristic. That is, $Q \circ \mathbb{L}^r$ is the zero polynomial over a field of characteristic $p \neq 2$. Further assuming that the coefficients of Q are in $\{-1, 0, 1\}$ naturally suggests a notion of a "p-modular" balanced family $\mathcal{A}$, where the number of graphs from $\mathcal{A}$ in every scope is a multiple of p. As in Corollary 6.4, a complexity lower bound on the members of $\mathcal{A}$ holds over any field of characteristic p (and thus also over $\mathbb{Z}$).

- A vanishing—over a finite field, say GF(2)—as a function (i.e., all evaluations are zero). In other words, $Q \circ \mathbb{L}^r = 0$ in the quotient ring $\mathrm{GF}(2)[\alpha]/({\alpha^i_{j,k}}^2 - \alpha^i_{j,k})_{i,j,k}$. Accordingly, one might generalize Definition 3.1 of comparable graphs, so that two equally-sized graphs would be considered comparable if they are simply spanned by the same vertices—without the requirement on their degrees (this follows by considering, in the proof of Lemma 6.3, equality of sets instead of multisets). While this would imply complexity lower bounds over any field of characteristic 2, lower bounds obtained using this redefinition may not be generalizable to arbitrary fields.

- A vanishing where r is extremely large. A simple counting argument we describe below implies the existence of a nonzero polynomial Q over any field $\mathbb{F}$ that evaluates to the zero polynomial on $\mathbb{L}^r$ for r as large as $\left\lceil \frac{n^{d-1}}{d} \right\rceil - 1$ (and in fact, using an additional observation, even as large as $\left\lceil \frac{n^d}{nd-d+1} \right\rceil - 1$), which is somewhat larger than the rank lower bound $\Omega(\frac{n^{d-1}}{d \log d})$ implied by Theorem 3.12. However, such a polynomial Q might not be multilinear (and its coefficients might not be in $\{-1, 0, 1\}$). In particular, if $\mathbb{F}$ is a finite field, Q might evaluate to zero on all inputs, and hence could not be used to deduce complexity lower bounds. Let us describe such a counting argument using the notion of algebraic dependence. By definition, the polynomials $(\mathbb{L}^r(\alpha)_e)_{e \in [n]^d}$ are *algebraically dependent* if and only if there exists a nonzero polynomial Q that evaluates on them to the zero polynomial (in algebraic-geometric terms, the algebraic set consisting of such vanishing polynomials is said to be *secant to a Segre variety*). It is known that if the number of polynomials $\mathbb{L}^r(\alpha)_e$ is larger than the number of their variables (that is; if $n^d > rdn$) then they are indeed algebraically dependent, and this holds for $r = \left\lceil \frac{n^{d-1}}{d} \right\rceil - 1$.

7 Results for more computational models

In this section we show that the definition of a balanced family can be naturally extended to imply complexity lower bounds with respect to various arithmetical computational models. We exemplify this by considering the case of homogeneous arithmetic circuits. An arithmetic circuit is said to be *homogeneous* if every intermediate polynomial it computes is homogeneous. It is well known (this is implicit in the work of Strassen [26]) that if a degree-d homogeneous polynomial has an arithmetic circuit of size s (i.e., that uses s arithmetical operations) then it has a homogeneous arithmetic circuit of size at most $O(d^2 s)$. Therefore, large lower bounds on the size of homogeneous arithmetic circuits can imply large lower bounds on the size of general arithmetic circuits.

In this section we take *d-graph* (or sometimes simply *graph*) to mean a (not necessarily partite) d-uniform hypergraph, i.e., where each edge includes precisely $d \geq 2$ vertices. All the graphs we consider are defined over the same vertex set $[n]$ for some positive integer n. The *adjacency polynomial* of a d-graph G is the polynomial

$$P_G(x_1, \ldots, x_n) = \sum_{e \in G} \prod_{v \in e} x_v.$$

Directed graphs. A *directed* edge is a pair (T, H), whose *tail* T and *head* H are sets of vertices. A directed graph is a collection of directed edges. A directed edge (T, H) is an *orientation* of an edge e if it partitions e, that is, $e = T \cup H$, and moreover, $|T|, |H| \leq \lfloor \frac{2}{3}|e| \rfloor$. An orientation of a graph is obtained by orienting each of its edges. For a directed graph $\vec{H}$, we denote H its underlying graph. We say that two directed graphs are *comparable* if they have the same multiset of tails, and similarly for heads (i.e., each tail that appears in one graph also appears as a tail in the other graph the same number of times, and similarly for heads).

Balanced families. Let $\mathcal{A} = \mathcal{A}^+ \cup \mathcal{A}^-$ be a family of d-graphs. We say that $\mathcal{A}$ is *t-balanced* if the following holds. For every graph G and every oriented t-subgraph[14] $\vec{H}$ of G, the number of directed graphs $\vec{H}'$ comparable to $\vec{H}$ such that $(G \setminus H) \cup H'$ lies in $\mathcal{A}^+$ is the same as for $\mathcal{A}^-$. It is not too hard to see that a family of d-partite hypergraphs that is t-balanced according to this definition is in particular t-balanced according to the definition in Section 3. The definition here is more restrictive since, intuitively, in order to obtain the rewirings of a t-subgraph, one cannot simply rearrange its vertices between the edges arbitrarily, but instead must only rearrange entire "blocks" as given by the orientation of each of the t edges. We next show that arithmetic circuit lower bounds can be obtained through balanced families.

Theorem 7.1. *Let F be a d-graph of cardinality m, and suppose that F is contained in a t-balanced family (as defined above). Then the size of any homogeneous arithmetic circuit computing P_F is $\Omega(m/t)$, over any field.*

We in fact prove a stronger result, which we describe next. Consider a representation of a degree-d homogeneous polynomial as a sum $\sum_i P_i Q_i$ where P_i, Q_i are homogeneous polynomials of degree at most $\lfloor \frac{2}{3}d \rfloor$. Let us call such a representation *systematic*. By a result of Raz [21], if a degree-d ($d \geq 2$) homogeneous polynomial p has a homogeneous arithmetic circuit of size s, then p has a systematic representation using $O(s)$ summands (see the proof of Proposition 2.7

[14] That is, every orientation of every t-subgraph.

in [21]).[15] Theorem 7.1 thus follows from the stronger theorem below.

Theorem 7.2. *In the notation of Theorem 7.1, the number of summands in any systematic representation of P_F is at least $\left\lceil \frac{m}{t-1} \right\rceil$, over any field.*

We sketch an algebraic proof along the lines of those in Section 6.

Proof. (Sketch) Let p be a multilinear homogeneous polynomial of degree d ($d \geq 2$) over a field $\mathbb{F}$, and consider a systematic representation $p = \sum_{i=1}^{s} P_i Q_i$ with s summands (without loss of generality, $\deg(P_i Q_i) = d$ for all i). For every $1 \leq i \leq s$, denote $(\alpha^i_T)_T, (\beta^i_H)_H \in \mathbb{F}$ the coefficients of the multilinear monomials in P_i and Q_i, respectively (we will only consider multilinear monomials in this proof). It follows that the coefficient of the (multilinear) monomial $\prod_{i \in e} x_i$ in $p = p(x_1, \ldots, x_n)$ is given by

$$\beta_e := \sum_{i=1}^{s} \sum_{T \,\dot\cup\, H = e} \alpha^i_T \beta^i_H,$$

where the innermost summation is over all orientations of e (in fact, over all orientations (T, H) of e where $|T| = \deg(P_i)$ and $|H| = \deg(Q_i)$).

Let $\mathcal{A} = \mathcal{A}^+ \,\dot\cup\, \mathcal{A}^-$ be a t-balanced family of size-m d-graphs. Let $Q_{\mathcal{A}} = Q_{\mathcal{A}}(X_e)$ be a (homogeneous, multilinear) polynomial over $\mathbb{F}$ whose variables correspond to edges of the complete d-graph (over $[n]$), and whose support corresponds to $\mathcal{A}$, namely,

$$Q_{\mathcal{A}} = \sum_{G \in \mathcal{A}^+} \prod_{e \in G} X_e - \sum_{G \in \mathcal{A}^-} \prod_{e \in G} X_e.$$

Viewing the β_e's as polynomials in the variables $\alpha := (\alpha^i_T, \beta^i_H)_{i \in [s], T, H}$, we consider the composed polynomial $Q' := Q_{\mathcal{A}}(\beta_e) \in \mathbb{F}[\alpha]$. Expanding any monomial $\prod_{e \in G} \beta_e$, it is easy to see that every term determines, for each edge of G, both a color (from $[s]$) and an orientation.

Fix $s < \left\lceil \frac{m}{t-1} \right\rceil$, so that every s-colored oriented graph contains at least t monochromatic (directed) edges. Using an argument similar to that at the end of the proof of Proposition 6.2, which, roughly speaking, reduces the analysis of "color-wise comparable" colored (directed) graphs to that of comparable (directed) graphs, as well as using an argument similar to that used in the proof of Proposition 4.6, the following can be shown. Consider any t-subgraph $\vec{H}$ of an oriented graph $\vec{G}$. Assume that the number of directed graphs $\vec{G'}$—obtained by replacing $\vec{H}$ with a comparable directed graph—whose underlying graph is from $\mathcal{A}^+$ is the same as from $\mathcal{A}^-$. Then each coefficient of Q' vanishes (for any $\mathbb{F}$). It is easy to see that this assumption is equivalent to $\mathcal{A}$ being t-balanced.

We conclude that $Q_{\mathcal{A}}$ vanishes on the coefficients β_e of all degree-d multilinear homogeneous polynomials that require less than $\left\lceil \frac{m}{t-1} \right\rceil$ summands for their systematic representation. Since $Q_{\mathcal{A}}$ does not vanish on any graph in its support (i.e., if $F \in \mathcal{A}$ then evaluating $Q_{\mathcal{A}}(X_e)$ by setting $X_e = 1$ if $e \in F$ and $X_e = 0$ otherwise—these are the coefficients of P_F—yields ± 1), the proof is complete. $\square$

[15] Note that the statement in [21] includes an extra factor due to homogenization and reduction to depth-4, which are not used here. We also do not use the fact that the P_i's and Q_i's can all be computed by an arithmetic circuit of size $O(s)$.

Notice that the above proof clearly applies even if we add arbitrary non-multilinear monomials to the adjacency polynomial of a balanceable graph.

One can also obtain a variant of the above theorem by considering depth-4 arithmetic circuits instead of general ones. Suppose that we modify the definition of an edge orientation so as to partition any edge into $O(\sqrt{d})$ blocks, each containing $O(\sqrt{d})$ vertices. Then it follows from the work of Koiran [14] (extending the work of Agrawal and Vinay [1]) that if we manage to obtain a lower bound of s on a d-graph (i.e., by putting it in a t-balanced family of size-m graphs where $\left\lceil \frac{m}{t-1} \right\rceil = s$), then a complexity lower bound of $s^{\Omega(1/\sqrt{d}\log d)}$ follows for general arithmetic circuits.

It is rather straightforward to generalize the above arguments to many other natural computational models, by appropriately modifying the definition of an edge orientation. For example, it is not hard to see that for (homogeneous) depth-3 arithmetic circuits, the corresponding notion of an edge orientation simply orders its vertices.

Remark. It is not hard to check that Theorem 3.12 can be extended to apply to balanced families as defined in this section. Namely, there exists a t-balanced family of size-m n-vertex d-graphs for $m = \Omega\big(\binom{n}{d}\big)$ and $t = O\big(\binom{n}{d'}\big)$ where $d' = \left\lfloor \frac{2}{3}d \right\rfloor$.

Acknowledgements. I thank Oded Regev for numerous invaluable comments, and for a careful reading of earlier drafts. I thank Michael Forbes for discussions on combinatorial interpretations related to tensors.

References

[1] M. Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 67–75, 2008.

[2] A. Alder. *Grenzrang und Grenzkomplexität aus algebraischer und topologischer Sicht*. PhD thesis, Universität Zürich, 1984.

[3] B. Alexeev, M. Forbes, and J. Tsimerman. Tensor rank: Some lower and upper bounds. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:10, 2011.

[4] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22(3):317–330, Feb. 1983.

[5] E. G. Belaga. Some problems involved in the computation of polynomials. *Dokl. Akad. Nauk SSSR*, 123:775–777, 1958.

[6] M. Bläser. A 5/2 n^2-lower bound for the rank of $n \times n$ matrix multiplication over arbitrary fields. In *40th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 45–50, 1999.

[7] R. A. Brualdi. Matrices of zeros and ones with fixed row and column sum vectors. *Linear Algebra and its Applications*, 33:159 – 231, 1980.

[8] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, 9(3):251–280, 1990.

[9] P. Erdős. On extremal problems of graphs and generalized graphs. *Israel Journal of Mathematics*, 2:183–190, 1964. 10.1007/BF02759942.

[10] B. Griesser. A lower bound for the border rank of a bilinear map. *Calcolo*, 23:105–114, 1986.

[11] D. Grigoriev. Some new bounds on tensor rank. *LOMI preprint E-2-1978*, 1978.

[12] V. Havel. A remark on the existence of finite graphs. *Casopis Pest. Mat.*, 80:477 – 480, 1955.

[13] J. JáJá. Optimal evaluation of pairs of bilinear forms. *SIAM J. Comput.*, 8(3):443–462, 1979.

[14] P. Koiran. Arithmetic circuits: the chasm at depth four gets wider. *CoRR*, abs/1006.4700, 2010.

[15] T. Lehmkuhl and T. Lickteig. On the order of approximation in approximative triadic decompositions of tensors. *Theor. Comput. Sci.*, 66(1):1–14, 1989.

[16] T. Lickteig. A note on border rank. *Inf. Process. Lett.*, 18(3):173–178, 1984.

[17] R. J. Lipton. Polynomials with 0-1 coefficients that are hard to evaluate. *SIAM J. Comput.*, 7(1):61–69, 1978.

[18] T. S. Motzkin. Evaluation of polynomials and evaluation of rational functions. *Bull. Amer. Math. Soc.*, 61:163, 1955.

[19] D. Mubayi. A hypergraph extension of Turán's theorem. *J. Comb. Theory, Ser. B*, 96(1):122–134, 2006.

[20] K. Mulmuley and M. A. Sohoni. Geometric complexity theory I: An approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.

[21] R. Raz. Elusive functions and lower bounds for arithmetic circuits. *Theory of Computing*, 6(1):135–177, 2010.

[22] R. Raz. Tensor-rank and lower bounds for arithmetic formulas. In *42nd ACM Symposium on Theory of Computing (STOC)*, pages 659–666, 2010.

[23] A. A. Razborov and S. Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997.

[24] H. J. Ryser. Combinatorial properties of matrices of zeros and ones. *Canad. J. Math.*, 9:371 – 377, 1957.

[25] A. Shpilka. Lower bounds for matrix product. *SIAM J. Comput.*, 32(5):1185–1200, 2003.

[26] V. Strassen. Vermeidung von divisionen. *J. Reine Angew. Math.*, 264:184–202, 1973.

[27] V. Strassen. Polynomials with rational coefficients which are hard to compute. *SIAM J. Comput.*, 3(2):128–149, 1974.

[28] V. Strassen. Rank and optimal computation of generic tensors. *Linear Algebra and its Applications*, 52/53:645–685, 1983.

[29] L. G. Valiant. Graph-theoretic arguments in low-level complexity. In *Proceedings of the 6th Symposium on Mathematical Foundations of Computer Science*, volume 53 of *LNCS*, pages 162–176. Springer, Sept. 1977.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de