Isomorphism of hypergraphs of low rank in moderately exponential time

László Babai[*] and Paolo Codenotti
University of Chicago
Email: {laci, paoloc}@cs.uchicago.edu

Abstract

We give an algorithm to decide isomorphism of hypergraphs of rank k in time $\exp\left(\widetilde{O}(k^2\sqrt{n})\right)$, where n is the number of vertices. (The rank is the maximum size of edges; the tilde refers to a polylogarithmic factor.) The case of bounded k answers a 24-year-old question and removes an obstacle to improving the worst case-bound for Graph Isomorphism testing. The best previously known bound, even for $k = 3$, was C^n (Luks 1999).

1 Introduction

NP problems are defined via a "witness space" $W(x)$ (where x is the input). $|W(x)|$ is typically exponential in $|x|$. Exhaustive search of $W(x)$ solves the problem in time $|W(x)| \operatorname{poly}(|x|)$. To reduce this time to $\exp\left((\ln |W(x)|)^{1-c}\right) \operatorname{poly}(|x|)$ for some constant $c > 0$ usually requires nontrivial insight; we call algorithms running within such a time bound "moderately exponential." Thus for the Graph Isomorphism problem, a moderately exponential algorithm would run in $\exp(O(n^{1-c}))$ where n is not the length of the input but the number of vertices (since the witness space in the definition consists of $n!$ permutations).

Moderately exponential algorithms for Graph Isomorphism were found in the wake of the introduction of group theoretic methods to the Graph Isomorphism problem ([Ba1, Lu1]); the best existing bounds are of the form $\exp(\widetilde{O}(\sqrt{n}))$ and arise from a combination of Luks's seminal work on divide-and-conquer methods to manipulate permutation groups [Lu1], combined with a combinatorial trick due to Zemlyachenko. The result appears in [BL] (1983); Luks subsequently refined the bound to $\exp\left(O(\sqrt{n \log n})\right)$ (see [BKL], 1984). This remains the state of the art after a quarter century; reducing the leading $\sqrt{n}$ term to $n^{1/2-c}$ in the exponent is a major open problem. For an important special case, that of strongly regular graphs, which admit a completely elementary $\exp\left(\widetilde{O}(\sqrt{n})\right)$ isomorphism test [Ba2], the exponent of the exponent was reduced to 1/3 by Spielman in 1996 [Sp].

While isomorphism of explicit hypergraphs is polynomial-time equivalent to Graph Isomorphism, reductions do not preserve moderate exponentiality, so the natural question arises whether isomorphism of hypergraphs, even of hypergraphs of bounded rank (bounded size of edges) can be tested in moderately exponential time. This question was stated in [BL] in 1984, and it was pointed out there that the absence of a moderately exponential bound for 4-uniform hypergraphs is an obstacle to improving the $\widetilde{O}(\sqrt{n})$ bound to $n^{1/2-c}$ in the exponent of the bound for Graph Isomorphism, adding to the significance of the problem.

In an important development, Luks reduced the trivial factorial bound to simply exponential (C^n) for testing isomorphism of hypergraphs of any rank [Lu3] (1999) (again, n is the number of vertices). However, this bound does not qualify as "moderately exponential," and Luks reiterates the long-standing open problem of moderately exponential isomorphism test for hypergraphs of bounded rank.

In the present paper we resolve this old problem, not only for bounded rank but even for rather large ranks.

Recall that a hypergraph $X = (V, E)$ consists of a set V of vertices and a set E of "edges;" the edges are subsets of V. The *rank* of X is the maximum size of its edges.

We state our main result. As above, we use the soft-Oh notation to suppress polylogarithmic factors, so for a function $f(n)$ we write $\widetilde{O}(f(n))$ to denote the class of functions $f(n)(\log n)^{O(1)}$. So for instance $n! = \exp(\widetilde{O}(n))$.

Theorem 1.1. *Isomorphism of hypergraphs of rank k with n vertices can be tested in* $\exp\left(\widetilde{O}(k^2\sqrt{n})\right)$.

Note that this bound is $\exp(\widetilde{O}(\sqrt{n}))$ for bounded k.

A key ingredient of our procedure is a subexponential algorithm for the Coset Intersection problem: given two subcosets of the symmetric group, determine their intersection.

[*] Partially supported by NSF Grant CCF-0830370. Any views expressed in this paper are the authors' and do not necessarily reflect the views of the NSF.

The intimate connection of this problem to Graph Isomorphism was first highlighted in Luks's seminal paper [Lu1]; an $n^{O(\sqrt{n})}$ algorithm was found by the first author:

Theorem 1.2. *([Ba7]) The Coset Intersection Problem can be solved in $n^{O(\sqrt{n})}$ time. In fact, it can be solved in time $n^{O(1)}m^{O(\sqrt{m})}$ time where m is the length of the longest orbit of the two groups.*

The $n^{O(\sqrt{n})}$ result was first described in [Ba5] (1983). It was included with a sketch of the proof in [BKL]. The proof of the main structural group theoretic lemma was given in full in [BBT]. While the outline of the coset intersection algorithm given in [BKL] omits key details, a full proof is now available [Ba7].

Other ingredients include the moderately exponential Graph Isomorphism test and a combinatorial lemma (Lemma 3.10).

2 Overall strategy

For the purposes of recursion, we shall need to consider the more general G-isomorphism problem, i. e., the problem of finding all isomorphisms between two objects X, Y on the same underlying set (vertex set) that belong to a given permutation group G. If $\mathrm{ISO}(X,Y)$ denotes the set of isomorphisms of X and Y then the set of their G-isomorphisms is $\mathrm{ISO}(G;X,Y) = G \cap \mathrm{ISO}(X,Y)$.

The overall scheme of our procedure is gradual approximation. In each round we consider an invariant. If the invariant fails to yield isomorphism rejection, we bring Y "closer" to X in the sense that X and Y now look identical with respect to the invariant. We reduce G in the process by making it respect the invariant. A simple illustration of this principle is to compare the degrees of the vertices of X and Y; if this comparison does not yield isomorphism rejection then we move each vertex of Y to a vertex of X of the same degree and reduce G to its subgroup that preserves the degree of each vertex of X.

We describe this idea on a slightly more formal level. While searching for irregularities in X, we may discover a proper subgroup $H \leq G$ such that $\mathrm{Aut}(G;X) \leq H$. Repeating the same process for Y either we refute G-isomorphism of X and Y or find a subgroup $K \leq G$ such that $\mathrm{Aut}(G;Y) \leq K$ and $K = \sigma H \sigma^{-1}$ is a conjugate of H for some $\sigma \in G$ which is computed along the way. Now we replace Y by Y^{σ} and G by H.

The goal is to reach a point when G becomes a subgroup of the automorphism group $\mathrm{Aut}(X)$ (so $\mathrm{Aut}(G;X) = G$); once this is the case, we conclude that either $X = Y$ or X and Y are not G-isomorphic.

Another aspect of our procedure is that it is recursive. If we find any irregularity in X, we either reject Y (if the same irregularity is not found in Y), or we split X as well as Y into "more regular" parts, determine the isomorphisms of the parts, and paste the results together.

An illustration of how this works is the case of invariant coloring of edges. (For instance, for graphs, we may color edges according to the number of triangles containing them.) Isomorphisms preserve this coloring. Now given the sets of isomorphisms of the corresponding pairs of edge-color-classes of the two hypergraphs, we need to take the intersection of these sets. Noting that the set of isomorphisms of X and Y is a coset, we see that this last step is an instance of Coset Intersection. This idea permits us to work with highly regular structures only.

Finally, following Luks's divide-and-conquer idea for permutation groups [Lu1], we need to delve into the group structure to some depth. As in [Ba7], the bottleneck arises from transitive groups with a large alternating or symmetric group action on a set of blocks of imprimitivity ("giant action"); handling this case constitutes the principal technical contribution of this paper. For the divide-and-conquer to succeed, we need to switch from hypergraphs to k-partite k-hypergraphs; this is analogous to moving to bipartite graphs from graphs.

3 Preliminaries

3.1 Groups

The *symmetric group* $\mathrm{Sym}(\Omega)$ acting on the set Ω consists of all permutations of Ω; the operation is composition. The *alternating group* $\mathrm{Alt}(\Omega)$ is the (unique) subgroup of index 2 in $\mathrm{Sym}(\Omega)$ consisting of the even permutations. Subgroups of $\mathrm{Sym}(\Omega)$ are called *permutation groups*; we refer to Ω as the *permutation domain*. The *degree* of a permutation group is the size of the permutation domain.

If $\Omega = [n] := \{1, \dots, n\}$ then we write S_n for $\mathrm{Sym}(\Omega)$ and A_n for $\mathrm{Alt}(\Omega)$.

Permutation groups will be represented by a list of generators; to "compute a subgroup" means to find generators for the subgroup.

Basic manipulation of permutation groups, including membership testing and finding the order, can be done in polynomial time ([Si1, Si2, FHL, Kn], see [Se]). It follows that from any list of generators, a minimal (non-redundant) list can be computed in polynomial time. Such a list has $\leq 2n$ elements, where $n = |\Omega|$ is the degree ([Ba6]).

For a group G we use the notation $H \leq G$ to express that H is a *subgroup*. If $H \leq G$ and $\sigma \in G$ then $H\sigma$ is a right coset of H in G. We shall call $H\sigma$ a *subcoset* of G. We also include the empty set among the subcosets (so this is $H\sigma$ for empty H, but the empty set is *not* a subgroup). Under this definition, intersection of subcosets is again a subcoset, namely, $H_1\sigma_1 \cap H_2\sigma_2$ is either empty or it is a coset of the subgroup $H_1 \cap H_2$. A coset $H\sigma$ is given by a set of generators of the group H and a representative $\tau \in H\sigma$.

3.2 Permutation groups: basic concepts

We review the basic concepts of the theory of permutation groups below; for a more detailed introduction we refer to [Wi, Lu1, Se].

An *action* of the group G on the set Ω is a homomorphism $\varphi : G \to \mathrm{Sym}(\Omega)$. The action is *faithful* if the kernel $\ker(\varphi)$ is trivial. We write the action in the exponent: for $\sigma \in G$ and $x \in \Omega$ we write $x \mapsto x^\sigma$. We suppress φ which will always be clear from the context. An action on Ω induces an action on the subsets of Ω: for $\Psi \subseteq \Omega$ we write $\Psi^\sigma = \{x^\sigma : x \in \Psi\}$.

For a permutation group $H \le \mathrm{Sym}(\Omega)$ we use the phrase "the action of G on Ω is H" if $\mathrm{Im}(\varphi) = H$.

Given an action $G \to \mathrm{Sym}(\Omega)$ and a set $\Psi \subseteq \Omega$, we define the *(pointwise) stabilizer* of Ψ in G to be $G_\Psi = \{\sigma \in G : (\forall x \in \Psi)(x^\sigma = x)\}$. Applying this notation to the induced action on subsets we obtain the notation $G_{\{\Psi\}} = \{\sigma \in G : \Psi^\sigma = \Psi\}$ for the *setwise stabilizer* of Ψ in G.

A subset $\Psi \subseteq \Omega$ is *G-invariant* if $(\forall \sigma \in G)(\Psi^\sigma = \Psi)$, i. e., $G_{\{\Psi\}} = G$. The *orbits* of the action are the equivalence classes of the relation on Ω defined as "$x \sim_G y$ if $(\exists \sigma \in G)(x^\sigma = y)$." In other words, the orbits are the minimal G-invariant subsets. The action is *transitive* if there is just one orbit, namely Ω. The orbit of $x \in \Omega$ is the set $x^G = \{x^\sigma : \sigma \in G\}$. We say that $|x^G|$ is the *length* of this orbit. Note that $|x^G| = |G : G_x|$ because the right coset $G_x\sigma$ consists of all $\tau \in G$ such that $x^\tau = x^\sigma$.

A *block of imprimitivity* is a nonempty subset $B \subseteq \Omega$ such that $(\forall \sigma \in G)(B^\sigma = B$ or $B^\sigma \cap B = \emptyset)$. A *partition* $\Omega = B_1 \dot\cup \ldots \dot\cup B_t$ into nonempty "blocks" B_i is G-invariant if $(\forall \sigma \in G)(\forall i \le t)(\exists j \le t)(B_i^\sigma = B_j)$. Such a partition is called a *system of imprimitivity* or a *block system.* The blocks of such a partition are blocks of imprimitivity. Conversely, a block of imprimitivity is a block in a unique invariant partition, assuming that the block has nonempty intersection with each orbit.

We say that the G-action on Ω is *primitive* if it is transitive, $|\Omega| \ge 2$, and Ω has no nontrivial G-invariant partition.

If $\mathfrak{B} = \{B_1, \ldots, B_t\}$ is a system of blocks then there is a natural action $G \to \mathrm{Sym}(\mathfrak{B})$. A *minimal block system* for a transitive action is a system $\mathfrak{B}$ of $t \ge 2$ blocks which has no nontrivial invariant coarsening. This is equivalent to saying that the G-action on $\mathfrak{B}$ is primitive.

The following fact will be central to our strategy.

Theorem 3.1. *Let $G \le S_n$ be a primitive group other than S_n or A_n. If n is sufficiently large then $|G| < n^{\sqrt{n}}$.*

This is a consequence of the classification of finite simple groups (cf. [Cam]). A slightly weaker $\exp(4\sqrt{n}\ln^2 n)$ upper bound, sufficient for our application, has an elementary combinatorial proof [Ba3, Ba4, Py]. We shall also need the following related but elementary estimate.

Fact 3.2. *Let $G = A_n$ or S_n and let $H < G$ be transitive but not A_n or S_n. If n is sufficiently large then $|G : H| \ge 2^n/\sqrt{2n}$.*

3.3 Reduction to k-partite k-hypergraphs

A *k-partite set* Ω is an ordered family of k sets $\Omega = (\Omega_1, \ldots, \Omega_k)$. We think of the Ω_i as being disjoint (replacing, as usual, Ω_i by $\Omega_i \times \{i\}$). Let $\mu(\Omega) = \dot\bigcup \Omega_i$ be the set of "points" in Ω, and let $\pi(\Omega) = \prod \Omega_i$. We shall think of the elements of $\pi(\Omega)$ as the *transversals* of Ω, i. e., those subsets of $\mu(\Omega)$ which intersect each Ω_i in exactly one element. We say that the k-partite set $\Psi = (\Psi_1, \ldots, \Psi_k)$ is a *subset* of Ω ($\Psi \subseteq \Omega$) if for each i, $\Psi_i \subseteq \Omega_i$.

If Ω and Ψ are k-partite sets then by a function $\sigma : \Omega \to \Psi$ we mean a function $\sigma : \mu(\Omega) \to \mu(\Psi)$ that respects the parts: $\Omega_i^\sigma \subseteq \Psi_i$. We say that a function $\sigma : \Omega \to \Psi$ is a bijection between Ω and Ψ if it is a bijection between $\mu(\Omega)$ and $\mu(\Psi)$.

For any $I \subseteq [k]$, let $\rho_I(\Omega) = (\Omega_i : i \in I)$, the $|I|$-partite set that is the restriction of Ω to I. We call the elements of $\pi(\rho_I(\Omega))$ the I-partial transversals or I-transversals of Ω. If $I \subseteq J \subseteq [k]$ and e is a J-partial transversal then the restriction of e to I is an I-partial transversal which we denote by $\mathrm{pr}_I(e)$. For a set E of J-partial transversals, we define the *projection* $\mathrm{pr}_I(E) = \{\mathrm{pr}_I(e) : e \in E\}$. Most of the time, we shall use these definitions with $J = [k]$.

Definition 3.3. *A* k-partite k-hypergraph $X = (V, E)$ *is defined by a k-partite set $V = (V_1, \ldots, V_k)$ of vertices and a set $E \subseteq \pi(V)$ of edges. We call the V_i the "parts" of the vertex set V. The* product-size *of X is $\Pi_0(X) = \prod_i |V_i|$.*

A *partial edge* is a subset of an edge. The partial edges are partial transversals; and *I-partial edge* is an I-partial transversal.

Let $X = (V; E)$ and $X' = (V'; E')$ be k-partite k-hypergraphs, where $V = (V_1, \ldots, V_k)$ and $V' = (V_1', \ldots, V_k')$. An isomorphisms of X and X' is a $V \to V'$ bijection which preserves edges, i. e., $e \in E$ iff $e^\sigma \in E'$. If this is the case, we write $X' = X^\sigma$. $\mathrm{ISO}(X, X')$ denotes the set of $X \to X'$ isomorphisms; and $\mathrm{Aut}(X) = \mathrm{ISO}(X, X)$ is the automorphism group of X.

Clearly, in discussing the isomorphism problem for X and X' we may restrict ourselves to the case when $V = V'$, i. e., $V_i = V_i'$ for all i.

We use the following notation: for a t-partite set $\Omega = (\Omega_1, \ldots, \Omega_t)$, let $\mathrm{SymPr}(\Omega)$ denote the group $\mathrm{Sym}(\Omega_1) \times \cdots \times \mathrm{Sym}(\Omega_t)$ in its action on $\mu(\Omega)$.

The following is immediate from Theorem 1.2.

Theorem 3.4. (Coset Intersection for k-partite sets) *Given a k-partite set $\Omega = (\Omega_1, \dots, \Omega_k)$, subgroups $G, H \le \mathrm{SymPr}(\Omega)$,and elements $\sigma, \tau \in \mathrm{SymPr}(\Omega)$, one can determine $G\sigma \cap H\tau$ in time $\mathrm{poly}(k) m^{O(\sqrt{m})}$, where $m = \max_i |\Omega_i|$.*

For the rest of this section, let $X = (V, E)$ and $Y = (V, F)$ be two k-partite k-hypergraphs over the same vertex set $V = (V_1, \dots, V_k)$. With this notation, the isomorphisms of X and Y form a subcoset of $\mathrm{SymPr}(V_1, \dots, V_k)$; our goal is to find this subcoset. We need to generalize this problem to finding all isomorphisms within a subcoset.

Definition 3.5. *Let $L \subseteq \mathrm{SymPr}(V_1, \dots, V_k)$ be a set of permutations. The set of L-isomorphisms between X and Y is $\mathrm{ISO}(L; X, Y) = L \cap \mathrm{ISO}(X, Y)$. The set of L-automorphisms of X is $\mathrm{Aut}(L; X) = \mathrm{ISO}(L; X, X) = L \cap \mathrm{Aut}(X)$.*

If L is a subcoset of $\mathrm{SymPr}(V_1, \dots, V_k)$ then so is $\mathrm{ISO}(L; X, Y)$. This is then the form in which we prove our main result:

Problem 3.6 (H-isomorphism of k-partite k-hypergraphs)**.** *Let X and Y be two k-partite k-hypergraphs over the same vertex set $V = (V_1, \dots, V_k)$. Let H be a subcoset of $\mathrm{SymPr}(V_1, \dots, V_k)$. Find the subcoset $\mathrm{ISO}(H; X, Y)$.*

Theorem 3.7. *Problem 3.6 can be solved in $\exp\left(\widetilde{O}(k^2\sqrt{n})\right)$ time, where $n = \max_{i=1}^k |V_i|$.*

Note that $\mathrm{ISO}(H\sigma; X, Y) = \mathrm{ISO}(H; X, Y^{\sigma^{-1}})$, for any set $H \subseteq \mathrm{SymPr}(V_1, \dots, V_k)$. So we may assume H is a group.

Observation 3.8. *For a coset H, H-isomorphism of hypergraphs of rank k with n vertices reduces to H^*-isomorphism of k-partite k-hypergraphs with parts of size at most n, and an increase of a factor of $\le k!$ in the number of edges. The k-partite k-hypergraphs and the coset H^* can be constructed in time, linear in the size of the output.*

Proof. Let Y_1 and Y_2 be hypergraphs of rank k over the same vertex set W, with edge sets F_1 and F_2, resp. We construct the k-partite k-hypergraphs $X_i = (W, \dots, W; E_i)$, $i = 1, 2$ where $E_i = \{(w_1, \dots, w_k) : \{w_1, \dots, w_j\} \in F_i$ for some $j \le k$ and $w_{j+1} = w_{j+2} = \dots = w_k \in \{w_1, \dots, w_j\}\}$ (so $|F_i| \le k!|E_i|$)). Let moreover $H^* = \{(\sigma, \dots, \sigma) : \sigma \in H\}$. Then the hypergraphs Y_1 and Y_2 are H-isomorphic if and only if the k-partite k-hypergraphs X_1 and X_2 are H^*-isomorphic. □

3.4 Linked actions

Throughout this section, $\Omega = (\Omega_1, \dots, \Omega_k)$ will be a k-partite set and $G \le \mathrm{SymPr}\,\mathrm{Sym}(\Omega_i)$; $G_i \le \mathrm{Sym}(\Omega_i)$ will be the action on Ω_i.

We say that the Ω_i are *linked* under the G-action (or "the domains of the G_i are linked") if $\mu(\Omega)$ admits a G-invariant partition into transversals $B_1, \dots, B_m \subseteq \mu(\Omega)$ (for each i, j, $|B_j \cap \Omega_i| = 1$). In this case, $|\Omega_1| = \dots = |\Omega_k|$; the transversals define bijections between the Ω_i.

Observation 3.9. *If Ω_i and Ω_j are linked for every pair $i < j \le k$ then all the Ω_i are linked. Moreover, the statement $S(i, j)$ that "Ω_i and Ω_j are linked" is an equivalence relation on $[k]$.*

Discovering links will be one of our main structural goals; the following lemma will serve us in this.

Lemma 3.10. *Let $Y = (\Omega_1, \Omega_2; E)$ be a biregular bipartite graph (regular on both parts) of positive density $\le 1/2$ (i. e., $1 \le |E| \le |\Omega_1||\Omega_2|/2$). Let $G \le \mathrm{SymPr}(\Omega_1, \Omega_2)$. (a) If $2\sqrt{n} \le |\Omega_i| \le n$ for $i = 1, 2$, and the restriction of $\mathrm{Aut}(G; Y)$ to Ω_1 contains $\mathrm{Alt}(\Omega_1)$ then the degree of each vertex in Ω_2 is 1. (b) If, in addition, the restriction of $\mathrm{Aut}(G; Y)$ to Ω_2 contains $\mathrm{Alt}(\Omega_2)$ then Ω_1 and Ω_2 are linked.*

Proof Let d be the degree of the vertices in Ω_2; assume $d \ge 2$. As $\mathrm{Alt}(\Omega_1)$ acts on Ω_1, every set of d vertices in Ω_1 must share a neighbor. There are $\binom{|\Omega_1|}{d} > |\Omega_2|$ such sets; therefore, by the pigeon hole principle, two of them must share a common neighbor. But such a common neighbor has degree greater than d, a contradiction, proving part (a). For part (b), we see that Y is regular of degree 1, so the edges of Y form the desired block system. □

If G is isomorphic to all the G_i, i. e., the G-action is faithful on each Ω_i, then we say that G acts *diagonally* on Ω. We shall use the following well-known fact.

Fact 3.11. *If each G_i is nonabelian simple then G is a direct product of diagonal actions.*

In other words, if all G_i are nonabelian simple then one can partition $[k]$ as $[k] = I_1 \dot\cup \dots \dot\cup I_t$ such that $G = D_1 \times \dots \times D_t$ where D_i acts diagonally on the $|I_i|$-partite set $(\Omega_j : j \in I_i)$ and it acts trivially on all the other Ω_j.

We shall say that the G_i corresponding to the same I_i are *linked*. We note that being linked is an equivalence relation.

Now we turn to the case of greatest importance to us: when each G_i is an alternating group.

Fact 3.12. *Assume $G_i = \mathrm{Alt}(\Omega_i)$, $|\Omega_i| \ge 5$ and $|\Omega_i| \ne 6$. If $G_1, \dots, G_k$ are linked then their domains are linked.*

This is a consequence of the fact that $\mathrm{Aut}(A_n) = S_n$ for all $n \ge 5, n \ne 6$; (cf. e. g., Lemma 2.1 in [BBT]). We need one more well-known fact about the alternating groups.

Fact 3.13. *In A_n, the subgroup of smallest index has index n if $n \ge 5$, the second smallest $\binom{n}{2}$ if $n \ge 9$.*

3.5 Giant threshold, giant action

Our algorithm will operate with a global constant n, to be thought of as the initial number of vertices of the hypergraphs. Throughout the algorithm, including its recursive calls, each "part" V_i of our k-partite k-hypergraphs will satisfy $|V_i| \leq n$.

We say that a permutation group $G \leq S_d$ is a *giant* if (a) G is either S_d or A_d; (b) $\gamma < d \leq n$ where $\gamma = \max\{7, 2\sqrt{n}\}$ is our *giant threshold.*

Definition 3.14. *The G-action on V is of* **giant type** *if the following holds. G acts transitively on each V_i and each V_i admits a G-invariant partition $\mathfrak{B}_i$; and the G-action on each $\mathfrak{B}_i$ (by permuting the blocks of $\mathfrak{B}_i$) is a giant.*

An application of Luks's divide-and-conquer strategy [Lu1] will show that within our target time bound, the only case we need to worry about is when G is of giant type.

Definition 3.15. *Assume the G-action on V is of giant type. If the action of G on $\mathfrak{B}_i$ is $\mathrm{Sym}(\mathfrak{B}_i)$ then let $G[i]$ be the subgroup of index 2 in G which acts as $\mathrm{Alt}(\mathfrak{B}_i)$ on $\mathfrak{B}_i$; otherwise set $G[i] = G$. We say that the G-action is of* **strong giant type** *if $(\forall i)(G[i]$ acts transitively on $V_i)$.*

Note that in the alternative, some V_i is split into two orbits of equal size under $G[i]$; these orbits are blocks of imprimitivity under G.

4 Color reductions, regularization

In this section we generalize the isomorphism problem to k-partite k-hypergraphs with colored vertices and colored edges; colors are preserved by isomorphisms by definition. We show how to get rid of the colors using recursion or group intersection. Since any observed irregularity provides us with an invariant coloring, a result of this homogenization is that we only need to deal with highly regular objects.

4.1 Vertex-color reduction

Let $f, f' : V \to \Sigma$ be two colorings of V. Let $X = (V, E, f)$ and $Y = (V, F, f')$ be two vertex-colored k-partite k-hypergraphs. Now $\mathrm{ISO}(X, Y) = H \cap \mathrm{ISO}((V, E), (V, F))$ where $H = \mathrm{ISO}((V, f), (V, f'))$ is the subcoset of color-preserving maps from the colored set (V, f) to the colored set (V, f').

It is trivial to compute H in polynomial time. Now if $G \cap H = \emptyset$ then X and Y are not G-isomorphic. Otherwise let $H = K\sigma$ where $K \leq \mathrm{SymPr}(V)$; then $\mathrm{ISO}(G; X, Y^{\sigma^{-1}}) = \mathrm{ISO}(G \cap K; X, Y^{\sigma^{-1}})$. Here $Y^{\sigma^{-1}} = (V, F^{\sigma^{-1}}, f)$, so we obtained the following (with $G^* = G \cap K$).

Reduction 4.1. *G-isomorphism of vertex-colored k-partite k-hypergraphs X and Y reduces to G^*-isomorphism of X' and Y', color-refined versions of X and Y, resp., where $G^* \leq G$; X' and Y' are colored identically; and this coloring is G^*-invariant. The cost of the reduction is a single application of coset intersection.*

We note that this reduction in particular permits the *elimination of isolated vertices.* In this case we shall reduce G to $G[i]$ and apply Luks's divide-and-conquer.

4.2 G transitive on each part

Suppose $V_i = V_i' \dot\cup V_i''$ where V_i', V_i'' are nonempty and G-invariant. We use Luks's divide-and-conquer idea to reduce this case to a case where V_i is reduced to V_i' and then a case where V_i is reduced to V_i'' (the latter to be solved within the output of the former). As a result we obtain:

Reduction 4.2. *The general case recursively reduces to the subcases when each part is a G-orbit. The product-size $\Pi_0(X) = \prod |V_i|$ is additive under this recursion.*

We note that this reduction does not work directly for hypergraphs (even for graphs); this is the main reason why we had to switch to k-partite hypergraphs.

4.3 Reduction to strong giant type

Now we use the other element of Luks's divide-and-conquer strategy to reduce to giant type. Let $\mathfrak{B}_i$ be a minimal block system on V_i. If the G-action on $\mathfrak{B}_i$ is not a giant, let K_i be the kernel of the G-action on $\mathfrak{B}_i$. Then we split G into cosets of K_i and solve the isomorphism problem within each coset. This reduces the problem to $|G : K_i|$ instances of K_i-isomorphism. Since each block of $\mathfrak{B}_i$ is K_i-invariant, we then use the previous section to reduce V_i successively to each block of $\mathfrak{B}_i$.

We use Theorem 3.1 to assess the cost of this recursion. While the multiplicative size of the problem was divided by $|\mathfrak{B}_i|$, the number of instances was multiplied by $|G : K_i| < |\mathfrak{B}_i|^{\sqrt{|\mathfrak{B}_i|}} < |\mathfrak{B}_i|^{\sqrt{n}} < |\mathfrak{B}_i|^{\gamma}$ if $|\mathfrak{B}_i| \geq \gamma$, where γ is our "giant threshold" (Sec. 3.5); and $|G : K_i| \leq |\mathfrak{B}_i|! < |\mathfrak{B}_i|^{\gamma}$ if $|\mathfrak{B}_i| < \gamma$. So in all cases, $|G : K_i| < |\mathfrak{B}_i|^{\gamma}$, giving a recurrence of the type $T(N) \leq b^{\gamma} T(N/b)$ (where $b = |\mathfrak{B}_i|$), resulting in a recursion tree with fewer than N^{γ} leaves, where N is the product-size $\Pi_(X)$. The leaves correspond to cases where G is either of giant type or trivial.

Finally if the G-action is of giant type but not of strong giant type, we find a V_i with an imprimitive action with 2 blocks (see end of Section 3.5) and proceed as above.

4.4 Edge-color reduction

Let $f : E \to \Sigma$, $f' : F \to \Sigma$ be two edge-colorings. Isomorphisms preserve edge-color by definition; and for any $L \subseteq \mathrm{SymPr}(V_1, \dots, V_k)$, $\mathrm{ISO}(L; (X, f), (Y, f')) = L \cap \mathrm{ISO}((X, f), (Y, f'))$.

Observation 4.3. *Let L be a subcoset of $\mathrm{SymPr}(V)$. Then we can find $\mathrm{ISO}(L; (X, f), (Y, f'))$ by solving $|\mathrm{Im}(f)|$ instances of Problem 3.6 and applying coset intersection.*

Indeed, if $\mathrm{Im}(f) \neq \mathrm{Im}(f')$ then (X, f) and (Y, f') are not isomorphic. Otherwise, for each color i, create two k-partite k-hypergraphs, X_i, and Y_i, over the vertex set $V = (V_1, \dots, V_k)$, with edge sets $E_i = \{e \in E | f(e) = i\}$, and $F_i = \{e \in F | f'(e) = i\}$, resp. We have $\mathrm{ISO}(L; (X, f), (Y, f')) = \bigcap_{i \in \mathrm{Im}(f)} \mathrm{ISO}(L; X_i, Y_i)$. □

Therefore if we find an invariant coloring of the edges of X and Y, we reduce Problem 3.6 to the color classes. The resulting hypergraphs will be increasingly "regular" ("irregularities" can be used for an invariant split).

In particular we obtain the following:

Reduction 4.4. *The general case reduces to the case when all edges of X and Y belong to a single G-orbit (on $\pi(V)$). The reduction is additive in terms of the number of edges. The cost is $m - 1$ applications of Coset Intersection, where m is the number of G-orbits into which the edges of X are divided.*

4.5 Regularization

Definition 4.5. *The k-partite k-hypergraph $X = (V, E)$ is* fully regular *if $\forall I \subseteq J \subseteq [k]$, $\forall e_1, e_2 \in \mathrm{pr}_I(E)$*

$$|\{e_3 \in \mathrm{pr}_J(E) : e_1 \subseteq e_3\}| = |\{e_3 \in \mathrm{pr}_J(E) : e_2 \subseteq e_3\}|.$$

Reduction 4.6. *The general case reduces to the case when X and Y are fully regular with the same parameters. The reduction is additive in terms of the number of edges. The cost is $m - 1$ applications of Coset Intersection, where m is the number of fully regular subsets into which the edges of X are divided.*

Indeed, if the hypergraph is not fully regular then then we label the I-partial edges by the number of J-partial edges containing them, and then color each edge by the label of its restriction to I.

4.6 Link-type

Assume G is of strong giant type. Then each V_i has a unique minimal block system $\mathfrak{B}_i$ on which G acts as a giant. Consider the G-action on the k-partite set $\mathfrak{B} = (\mathfrak{B}_1, \dots, \mathfrak{B}_k)$. We write $i \sim_b j$ and say that the parts V_i and V_j are **block-linked** if the G-actions on $\mathfrak{B}_i$ and $\mathfrak{B}_j$ are linked. In this case we also say that v_i and v_j are block-linked for all $v_i \in V_i$ and $v_j \in V_j$.

We write $i \sim_e j$ and say that the parts V_i and V_j are edge-linked under edge $e = (v_1, \dots, v_k) \in E$ if the block containing $v_i \in V_i$ and the block containing $v_j \in V_j$ are linked. We call the equivalence relation $\sim_e$ the **link-type** of the edge e; it is a refinement of the equivalence relation $\sim_b$. We note that G preserves link-type; therefore, after Reduction 4.4, *all edges have the same link-type* and we may refer to $\sim_e$ without specifying e.

The link-type of I-partial edges is the restriction of $\sim_e$ to $I \subseteq [k]$.

Let $\mathfrak{J}$ denote the set of $\sim_b$-equivalence classes, and $\mathfrak{H}$ the set of $\sim_e$-equivalence classes (each of these partition $[k]$). We call the elements of $\mathfrak{J}$ **block-linked sections** and the elements of $\mathfrak{H}$ **edge-linked sections**.

Some further notation. For an edge-linked section h, let $M_h^1, \dots, M_h^{k(h)}$ be a G-invariant partition of $\{\mathfrak{B}_i : i \in h\}$ into transversals. We call the M_h^j **macro-blocks**. Note that each macro-block consists of one block from each part of an edge-linked section. Note further that all the edges in an edge-linked section will be within the macro-blocks. Let $\mathfrak{M}_h = \{M_h^s : s \in [k(h)]\}$ be the set of macro-blocks of edge-linked section h.

For any $v \in V$, we call $\mathrm{block}(v)$ the block containing v, and $\mathrm{macroblock}(v)$ the macro-block containing v.

4.7 The ultimate structure

Assume G is of strong giant type, all edges of $X = (V, E)$ belong to the same G-orbit, and X has no isolated vertices. We use the notation of Sec. 4.6. Let $f = (v_1, \dots, v_k) \in \pi(V)$ be a transversal ($v_i \in V_i$). We say that $f \in \hat{E}$ if (a) for each edge-linked section $h \in \mathfrak{H}$, there is an edge $e \in E$ such that $e|_h = f|_h$; and (b) if $i \not\sim_e j$ then v_i and v_j are not block-linked. Obviously $E \subseteq \hat{E}$.

We say that the pair (X, G) has the **ultimate structure** if the assumptions of the first sentence hold and $\hat{E} = E$.

Lemma 4.7. *Assume both (X, G) and (Y, G) have the ultimate structure. Then the G-isomorphisms of X and Y can be determined in time $O(n^{k\sqrt{n}})$.*

Proof. Let $\hat{G}$ be the largest supergroup of G in $\mathrm{SymPr}(V)$ which has the same blocks and the same block-link relation $\sim_b$. It is trivial to determine $\hat{G}$ in polynomial time. It suffices to find all $\hat{G}$-isomorphisms and then apply coset intersection with G. The $\hat{G}$-isomorphisms can be pieced together from the $\hat{G}$-isomorphisms of the restrictions of X and Y to each $\sim_e$ class. But each connected component of such a restriction is contained in a macro-block, so the pairwise isomorphisms of the connected components can be computed by brute force in $((\sqrt{n})!)^k$. □

5 Breaking symmetry

5.1 Individualization

To *individualize* a vertex $x \in V_i$ in X means to assign x a special color, say "red," in X, resulting in the structure X_x. For each $y \in V_i$, let $\sigma_y \in G$ move x to y. Now we have $\mathrm{ISO}(G, X, Y) = \bigcup_{y \in V_i} \mathrm{ISO}(G_x \, \sigma_y; X_x, Y_y)$.

This recurrence thus reduces an instance of the G-isomorphism to $\leq n$ instances of G_x-isomorphism. To individualize a set R of vertices means to individualize each $x \in R$ in succession; the cost is a factor of $\leq n^{|R|}$. By breaking the symmetry we hope to make a sufficient reduction in our "potential functions" to offset this factor.

5.2 Split

Let G act on V, with blocks of imprimitivity $\mathfrak{B}_i$ on V_i. Let $R \subseteq \mu(V)$, and $S \subseteq E$ or $S \subseteq V_i$ or $S \subseteq \mathfrak{B}_i$, such that S is invariant under the pointwise stabilizer $\mathrm{Aut}(G; X)_R$.

Definition 5.1. *Let $S' \subseteq S$, $S' \neq \emptyset$, $S' \neq S$. We say that S' is a* split *of S relative to R if S' is invariant under* $\mathrm{Aut}(G; X)_R$. *If $R = \emptyset$, we say that this is an* absolute split. *If $|S|/\gamma \leq |S'| \leq (1 - 1/\gamma)|S|$ and $|R| \leq 2k$ then we say this is a* large relative split*; otherwise it is a* small split.

Note that if we have an absolute split, then we can reduce one of the V_i or the $\mathfrak{B}_i$ or E. Therefore whenever we find an absolute split, we can recurse. If we find a large split relative to a "small" set R, we will individualize R pointwise, and recurse. Our R will always be the union of at most two partial transversals, hence the inequality $|R| \leq 2k$.

5.3 Reduction of G

Most steps of the algorithm involve reducing G. The new group H is always a quotient of a subgroup of the original; in particular, $|H| \leq |G|$. If this reduction is achieved through individualization of a subset $R \subseteq \mu(V)$ then we say we have a *large relative reduction of* G if (a) $|H| < |G| \exp(-\gamma/2)$; and (b) $|R| \leq 2k$.

This reduction will be used in the context of Fact 3.13.

6 Measures of progress

6.1 Events of progress

We now define a list of structures which will allow us *to make progress*. If we discover one of these structures, we take the recursive action stated in parentheses next to each structure. Henceforth, if we say "we make progress," this phrase will always refer to this list.

1. $|V_i| = 1$ for some i (reduction: $k \leftarrow k - 1$);
2. $\mathrm{Aut}(X)$-invariant vertex coloring (this includes the case when the action of G is intransitive on any of the V_i) (reduce to color-homogeneous parts, Sec. 4.1);
3. $\mathrm{Aut}(X)$-invariant edge-coloring (reduce to color-homogeneous parts, Sec. 4.4; this includes Reduction 4.4 (all edges belong to a single G-orbit));
4. The G-action on a set of maximal blocks in one of the V_i is not a giant (divide-and-conquer, Sec. 4.3);
5. Absolute split or large relative split (reduce to the two parts of the split);
6. Large relative reduction of G (reduce G).

The following result describes the nature of our algorithm; its proof follows in Section 7.

Theorem 6.1. *Our algorithm proceeds in a sequence of recursive phases each of which makes progress. The cost of each reduction is a single composite step.*

6.2 Potential; timing analysis

We introduce three "potential functions" on the pairs (X, G):

$$\Pi_1(X) = |E|\Pi_0(X) = |E| \cdot \prod_i |V_i|,$$

$$\Pi_2(G) = \exp \frac{4k \ln |G| \ln n}{\gamma},$$

$$\Pi(G, X) = \Pi_1(X)^{\gamma(2k+1)\ln n} \Pi_2(G).$$

The following theorem justifies the claimed running time for our algorithm. In the Theorem we consider the following operations as **elementary steps:** computing the coset of isomorphisms of a pair of graphs with $O(n)$ vertices; computing coset intersection of permutation groups acting on a k-partite domain in which each part has size $O(n)$; and operations requiring time $n^{O(k)}$. A **composite step** consists of any combination of $n^{O(k)}$ elementary steps.

Theorem 6.2. *Let $T(G, X)$ denote the number of composite steps required by the algorithm. Then*

$$T(G, X) \leq \Pi(G, X).$$

We show that this bound follows from Theorem 6.1. We set $\gamma' = \gamma(2k + 1) \ln n$.

Proof. By induction on $\Pi(G, X)$. The base cases are $k \leq 2$ (a single application of Graph Isomorphism) and the "ultimate structure" (Sec. 4.7) which is settled by Lemma 4.7.

We note that neither Π_1 nor Π_2 increases during recursive calls. We distinguish cases according to the last event of progress; the numbering follows that of Section 6.1.

1. The reduction $k \leftarrow k-1$ takes care of the unit cost.

2-3. Invariant coloring. We reduce the problem (X,G) to instances (X_i,G_i). Then $(\forall i)(\Pi_2(G) \leq \Pi_2(G_i))$, while
$$\Pi_1(X) = \sum_i \Pi_1(X_i). \qquad (1)$$
Moreover, $T(G,X) \leq 1 + \sum_i T(G_i,X_i)$. (The added unit covers the cost of the reduction.) By the inductive hypothesis $T(G_i,X_i) \leq \Pi(G,X)$. So, by (1),
$$\frac{T(G,X)}{\Pi_2(G)} \leq 1 + \sum_i \Pi_1(X_i)^{\gamma'} < \Pi_1(X)^{\gamma'}.$$

4. Divide-and-conquer for imprimitive groups. Following the argument outlined in Sec. 4.3, let b be the number of blocks; so we reduce to at most b^γ instances (X_i,G_i) where $(\forall i)(\Pi_1(X_i) = (1/b)\Pi_1(X))$; hence
$$\frac{T(G,X)}{\Pi_2(G)} < 1 + \sum_i \Pi_1(X_i)^{\gamma'}$$
$$\leq 1 + b^\gamma \left((1/b)\Pi_1(X)\right)^{\gamma'} < \Pi_1(X)^{\gamma'}.$$

5. Split. We reduce the instance (G,X) to two instances $(G_1,X_1),(G_2,X_2)$. Again $\Pi_2(G_i) \leq \Pi_2(G)$, and $\Pi_1(X) = \Pi_1(X_1) + \Pi_1(X_2)$. If the split is absolute then $T(G,X) \leq 1 + \sum_i T(G_i,X_i)$. By the inductive hypothesis, $T(G_i,X_i)/\Pi_2(G) < 1 + \sum_i \Pi_i(X_i)^{\gamma'} < \Pi_1(X)^{\gamma'}$. If the split is a large relative split then we have $\Pi_1(X_i) < (1-1/\gamma)\Pi_1(X)$ but we incur a $\leq n^{2k}$ multiplicative cost factor for individualization. So $T(G,X)/\Pi_2(G) < 2n^{2k}\left((1-1/\gamma)\Pi_1(X)\right)^{\gamma'} < \Pi_1(X)^{\gamma'}$. (Note that in all other events of progress, we could have set $\gamma' = \gamma$; here we needed the factor γ'/γ to counter the multiplicative cost of individualization.)

6. Large relative reduction of G. In this case, $\Pi_2(G)$ is reduced by a factor of $\leq n^{-2k}$, offsetting the $\leq n^{2k}$ multiplicative cost of individualization. □

It follows that our algorithm has the claimed running time since initially $\Pi_1(X) \leq n^{2k}$ and $|G| \leq (n!)^k < n^{nk}$.

7 The algorithm

We focus on $X = (V,E)$; if Y is found to differ, we reject isomorphism. Our goal is to keep making progress (or reject isomorphism) until the ultimate structure (Sec. 4.7) is reached. By Sec. 4, through a sequence of "progress events" (Sec. 6.1) we reach a state where (1) G is of strong giant type; (2) all edges lie in the same G-orbit; (3) X is fully regular and has no isolated vertices.

Notation 7.1. *Let $t \subseteq \mu(V)$. Let Z be either $\mathfrak{B}_i$ or V_i or $\mathfrak{M}_r$. We define $Z(t)$ as the subset of Z consisting of the elements not block-linked to any vertex in t.*

7.1 Vertex splitting: the tree argument

Let $T \subseteq \mathcal{P}(\mu(V))$ ($\mathcal{P}$ denotes the power-set). We say that an assignment $f : T \to \mathcal{P}(\mu(V))$ is *equivariant* if $f(t^\sigma) = f(t)^\sigma$ for every $t \in T$, $\sigma \in \mathrm{Aut}(G;X)$.

The following theorem is our key tool that will allow us to make progress any time we find a split. For $\mathfrak{L} \subseteq \mathfrak{H}$ we use $\mathrm{pr}_{\mathfrak{L}}(E)$ to denote the projection of E to $\bigcup \mathfrak{L}$.

Theorem 7.2 (tree argument)**.** *Let $\mathfrak{L} \subseteq \mathfrak{H}$, $i \in [k]$, and let $\{S_t \subseteq V_i | t \in \mathrm{pr}_{\mathfrak{L}}(E)\}$ be an equivariant assignment of subsets of V_i to elements of $\mathrm{pr}_{\mathfrak{L}}(E)$. If $\exists t \in \mathrm{pr}_{\mathfrak{L}}(E)$ such that S_t is a split of $V_i(t)$ then in a single composite step we can make progress.*

Proof. We will prove this by reverse induction on $|\mathfrak{L}|$. If $|\mathfrak{L}| = 0$, then we have an absolute split, hence progress.

If $|S_t| \neq |S_{t'}|$ for some $t,t' \in \mathrm{pr}_{\mathfrak{L}}(E)$, then we color $\mathrm{pr}_{\mathfrak{L}}(E)$ by $|S_t|$, which is progress. W.l.o.g. $|S_t| \leq |V_i|/2$. If S_t is a large split of V_i relative to t, progress. Otherwise not all blocks $\mathfrak{B}_i(t)$ contain elements of S_t. Call the blocks in $\mathfrak{B}_i(t)$ that contain any elements of S_t special, and let the set of macro blocks containing a special block be $\mathrm{sp}(t)$. Let $\ell \in \mathfrak{H}$ such that $i \in \ell$. If for any t, $\mathrm{sp}(t)$ is a large split of $\mathfrak{M}_\ell(t)$, then we have a large split of V_i, progress. Else, pick any $h = \{h_1,\dots,h_\alpha\} \in \mathfrak{L}$. Let $\mathrm{sp}(\alpha+1;t) = \mathrm{sp}(t)$. For every $1 < \beta \leq \alpha$, and every $u \in \mathrm{pr}_{\mathfrak{L}\setminus\{h_\beta,\dots,h_\alpha\}}(E)$, let $V(u,\beta) = \{v \in V_{h_\beta} : (u,v) \in \mathrm{pr}_{\mathfrak{L}\setminus\{h_{\beta+1},\dots,h_\alpha\}}\}$, and let $\mathrm{sp}(\beta;u) = \bigcup_{v \in V(u,\beta)} \mathrm{sp}(\beta+1,(u,v))$.

If any of these sets is a large split of $\mathfrak{M}_\ell$, progress. Else, we claim that $|\mathrm{sp}(\beta;u)| \leq |\mathfrak{M}_\ell|/\gamma = |\mathfrak{B}_i|/\gamma$, where $\gamma = 2\sqrt{n}$ is our "giant threshold." This follows by reverse induction on β. Indeed, by induction, $\mathrm{sp}(\beta;u)$ is the union of $\leq n/\gamma$ sets, each of size $\leq |\mathfrak{M}_\ell|/\gamma$, so $|\mathrm{sp}(\beta;u)| \leq |\mathfrak{M}_\ell|/2$ and therefore $|\mathrm{sp}(\beta;u)| \leq |\mathfrak{M}_\ell|/\gamma$ (otherwise it would be a large split).

For each $M \in \mathfrak{M}_h$ and $u \in \mathrm{pr}_{\mathfrak{L}\setminus\{h\}}(E)$, let $\mathrm{sp}(u,M) = \bigcup_{v \in M \cap V_{h_1}} \mathrm{sp}(1,(u,v))$. If any of these special sets are large splits of $\mathfrak{M}_\ell$, progress. Else we create graphs $Y(u)$ for each $u \in \mathrm{pr}_{\mathfrak{L}\setminus\{h\}}(E)$. If $h \neq \ell$ then $Y(u)$ will be bipartite with vertex set $(\mathfrak{M}_h(u),\mathfrak{M}_\ell(u))$; if $h = \ell$ then $Y(u)$ is a graph with vertex set $\mathfrak{M}_\ell(u)$. In both cases, the edge set is $\{(M \in \mathfrak{M}_h(u), N \in \mathfrak{M}_\ell(u)) : N \in \text{special}(u,M)\}$. We will refer to $\mathfrak{M}_h(u)$ and $\mathfrak{M}_\ell(u)$ as the parts of $Y(u)$.

Now for each $u \in \mathrm{pr}_{\mathfrak{L}\setminus\{h\}}$, let S_u be a largest orbit of $\mathrm{Aut}(Y(u))$ on V_j. If any of the S_u is a large split, progress. Else, if any S_u is a split, we are done by induction on $|\mathfrak{L}|$. Else, $\mathrm{Aut}(Y(u))$ is transitive on each part of $Y(u)$.

If the action of $\mathrm{Aut}(Y(u)$ is transitive on each part of $Y(u)$ but the not giant on either part, then we obtain a large reduction of G relative to u (Fact 3.13), progress. Else, we claim that $h \neq \ell$ and $Y(u)$ is a matching. This follows by Lemma 3.10 from the fact that the density of $Y(u)$ is positive and $\leq 1/2$. This means we found a linking between previously unlinked blocks, yielding a large relative reduction in G, progress.

To analyze the cost note that for each partial edge, we do polynomial work and one Graph Isomorphism test. The set of partial edges considered forms a tree, which is where the Theorem gets its name. $\square$

Theorem 7.3 (double tree argument)**.** *Let $\mathfrak{L}_j \subseteq \mathfrak{H}$ for $j = 1, 2$; let $i \in [k]$; and let $\{S_T | T \in \mathrm{pr}_{\mathfrak{L}_1}(E) \times \mathrm{pr}_{\mathfrak{L}_2}(E)\}$ be an equivariant assignment of subsets of $V_i(T)$ to elements of $\mathrm{pr}_{\mathfrak{L}_1}(E) \times \mathrm{pr}_{\mathfrak{L}_2}(E)$. If $\exists T \in \mathrm{pr}_{\mathfrak{L}_1}(E) \times \mathrm{pr}_{\mathfrak{L}_2}(E)$ such that S_T is a split of $V_i(T)$ then in a single composite step we can make progress.*

Proof. Fix $t_1 \in \mathrm{pr}_{\mathfrak{L}_1}(E)$ and perform the tree argument with respect to $\mathrm{pr}_{\mathfrak{L}_1}(E)$. If the tree argument makes progress that involves individualization, then we will individualize all of t_1 and make the same progress. Otherwise the tree argument makes progress by finding a small absolute split; but this is a split relative to t_1, therefore we need to invoke the tree argument again. $\square$

7.2 Structure

Theorem 7.4. *If we do not observe the ultimate structure then we make progress.*

Proof. For $h \in \mathfrak{H}$, let $\overline{h} = \mathfrak{H} \setminus \{h\}$. Fix any $h \in \mathfrak{H}$. For each $t \in \mathrm{pr}_{\overline{h}}(E)$ and $M \in \mathfrak{M}_h$ we define a hypergraph $H(t, M) = (M, E(t, M))$ where $E(t, M) = \{u \in \mathrm{pr}_h(E) \cup \pi(M) : (t, u) \in E\}$.

If $M \notin \mathfrak{M}_h(t)$ then $E(t, M) = \emptyset$ since E consists only of edges of the same link-type.

If there exist pairs (t, M), (t', M') such that $E(t, M), E(t', M') \neq \emptyset$ but $H(t, M) \ncong H(t', M')$, then the following is an invariant coloring of E. For any $e \in E$, let $t(e) = \mathrm{pr}_{\overline{h}}(e)$, and let $M(e) = \mathrm{macroblock}(\mathrm{pr}_h(e))$ be the macro-block containing $\mathrm{pr}_h(e)$. Then color e by the isomorphism type of $H(t(e), M(e))$.

If $E(t, M) = \emptyset$ for some $M \in \mathfrak{M}_h(t)$, then for every $s \in \mathrm{pr}_{\overline{h}}(E)$, the set of macro-blocks with empty hypergraphs is an invariant partition of $\mathfrak{M}_h(s)$. Moreover, for t the set of empty macro blocks is a split of $\mathfrak{M}_h(t)$, since we are assuming $\exists M$ such that $E(t, M) = \emptyset$, but $\exists M'$ such that $E(t, M') \neq \emptyset$ since t is a partial edge. Thus we make progress via the tree argument (Theorem 7.2).

Now if X does not have the structure, we can find some h^*, $t \in \mathrm{pr}_{\overline{h^*}}(E)$, and $u \in \mathrm{pr}_{h^*}(E)$ such that $(t, u) \notin E$. But since $u \in \mathrm{pr}_{h^*}(E)$, $\exists s \in \mathrm{pr}_{\overline{h^*}}(E)$ such that $(s, u) \in E$. Given this we will make progress. Let $r = h^*$.

For $s, t \in \mathrm{pr}_{\overline{r}}(E)$, and $M \in \mathfrak{M}_r$, let $H(s, t; M)$ be the superposition of the hypergraphs $H(s, M)$ and $H(t, M)$, that is, the colored hypergraph with the edges of $H(s, M)$ of one color and the edges of $H(t, M)$ of another color. Note that if X has the structure, then $H(s, t; M)$ will be the superposition of two identical hypergraphs. We can compute isomorphisms of the $H(s, t; M)$ by exhaustive search in time $\exp(O(k\sqrt{n}))$. We want to show that all the $H(s, t, M)$ are isomorphic, and that the automorphism groups $A(s, t, M) = \mathrm{Aut}(H(s, t; M))$ act transitively on all partial edges in M.

Let $\mathcal{H}$ be the set of isomorphism types of the $|r|$-partite $|r|$-hypergraphs $H(s, t, M)$ ($s, t \in \mathrm{pr}_{\overline{r}}(E)$, $M \in \mathfrak{M}_r$). Fix $H \in \mathcal{H}$. Let $S(s, t) = \{M \in \mathfrak{M}_r : H(s, t; M) \cong H\}$. This is an equivariant map on the pairs (s, t). If $S(s, t)$ splits $\mathfrak{M}_r$ for some s, t, H then we have a split of V_j ($j \in r$) and we make progress by the double tree argument (Thm. 7.3).

Else we conclude that the isomorphism type of $H(s, t, M)$ does not depend on M. Let us again fix $H \in \mathcal{H}$. Let $S(t) = \{s \in \mathrm{pr}_{\overline{r}}(E) : H(s, t; M) \cong H\}$. This again is an equivariant map. Again, if for some t, H this is a split, we make progress by the tree argument (Thm 7.2). Else we conclude that the isomorphism type of $H(s, t, M)$ does not depend on s; by symmetry, it also does not depend on t.

Let now $A(s, t, M) = \mathrm{Aut}(H(s, t, M))$. If $A(s, t, M)$ is not vertex transitive on the vertices of one of the parts of M then we have a split into orbits relative to s, t, and M. It is the same size split in all the macro blocks in $\mathfrak{M}_r$, so it is a large split since there are $\geq \gamma$ macro blocks. Progress.

Our next goal is to achieve that $A(s, t, M)$ is transitive on $E(s, t, M)$.

We iterate through r as follows. Let $r = \{i_1, \dots, i_{|r|}\}$ and let $r_\ell = \{i_1, \dots, i_\ell\}$. We prove by induction on ℓ that either $A(t, s; M)$ acts transitively on $\mathrm{pr}_{r_\ell}(H(t, s; M))$ or we make progress. We have shown this for $\ell = 1$ (vertex-transitivity). Let $\ell \geq 2$ be the smallest value such that $A(t, s; M)$ acts intransitively on $\mathrm{pr}_{r_\ell}(H(t, s; M))$. Then for any partial edge $e \in \mathrm{pr}_{r_{\ell-1}}(H(t, s; M))$, we get a split of $\mathrm{pr}_{r_\ell}(H(t, s; M))$ into orbits, and hence a split of V_{i_ℓ} relative to t, s and e. The split is large, since it is a split in every block of V_{i_ℓ}, progress.

We know that there exists a macro block $M \in \mathfrak{M}_r$ and pair of partial edges $s, t \in \mathrm{pr}_{\overline{r}}(E)$ such that $H(s, M) \neq H(t, M)$ (since we do not observe the ultimate structure). Look at the smallest ℓ (over all choices of $s, t \in \mathrm{pr}_{\overline{r}}(E)$, and $M \in \mathfrak{M}_r$) such that $H(s, M)$ and $H(t, M)$ are not identical when restricted to the first ℓ parts of M.

Note that ℓ cannot be 1: $A(s, t, M)$ acts transitively on the vertices in each part of M.

Note further that because $A(s, t, M)$ acts transitively both on $E(t, M)$ and on $E(s, M)$, these sets are either identical or disjoint, and so are any of their projections.

Define equivalence classes on $\mathrm{pr}_{\overline{r}}(E)$, where two partial edges s and t are equivalent if and only if $H(s, M)$ and $H(t, M)$ are identical up to level ℓ. We individualize one of the equivalence classes; this splits E. We claim this is a large (relative) split.

Indeed, fix a partial edge e of length $\ell - 1$ in $H(s, M)$. (Thus is also a partial edge of $H(t, M)$ for all t.) Let B be the block of M in V_{i_ℓ}. Let $R(t) = \{v \in B$ s.t. (t, e, v) is a partial edge$\}$. The $R(t)$ partition B. On the other hand, the relation $v \in R(t)$ defines a biregular bipartite graph on $(\mathrm{pr}_{\overline{T}}(E), B)$ (because of full regularity). It follows that all equivalence classes are the same size. Their number is at most $|B| \leq n/\gamma$. $\square$

8 Open Problems

The big open problem continues to be to reduce the $\widetilde{O}(\sqrt{n})$ term in the exponent of the complexity of Graph Isomorphism to $n^{1/2-\epsilon}$. Our result for 4-hypergraphs may open the path to a reduction of the exponent to $\widetilde{O}\left(n^{1/4}\right)$.

Another problem is to extend our moderately exponential isomorphism test to hypergraphs of rank $n^{1-\epsilon}$.

Finally, our work brings new life to an old dilemma: *isomorphism testing vs. canonical forms* (cf. [BL, FSS]). Because of the use of Coset Intersection, our algorithm does not yield canonical forms. As Gene Luks points out, his simply exponential (C^n) isomorphism test for hypergraphs (of arbitrary rank) [Lu3] does not yield canonical forms either. So the problem is, find canonical forms (a) for hypergraphs in simply exponential time and (b) for hypergraphs of bounded rank in moderately exponential time.

References

[Ba1] L. BABAI: Monte Carlo algorithms in graph isomorphism testing. Université de Montréal Tech. Rep. DMS 79-10, 1979 (pp. 42) http://people.cs.uchicago.edu/~laci/lasvegas79.pdf

[Ba2] L. BABAI: On the complexity of canonical labelling of strongly regular graphs. *SIAM J. on Computing* **9** (1980), 212-216.

[Ba3] L. BABAI: On the order of uniprimitive permutation groups. *Ann. Math.* **113** (1981), 553–568.

[Ba4] L. BABAI: On the order of doubly transitive permutation groups. *Inventiones Math.* **65** (1982), 473–484.

[Ba5] L. BABAI: *Permutation Groups, Coherent Configurations and Graph Isomorphism.* D.Sc. Thesis (Hungarian), Hung. Acad. Sci., April 1983.

[Ba6] L. BABAI: On the length of subgroup chains in the symmetric group. *Communications in Algebra* **14** (1986), 1729–1736.

[Ba7] L. BABAI: Coset intersection in moderately exponential time. *Chicago Journal of Theoretical Computer Science*, to appear. http://cjtcs.cs.uchicago.edu

[BBT] L. BABAI, R. BEALS, P. TAKÁCSI-NAGY: Symmetry and complexity. In: *Proc. 24th STOC*, ACM Press 1992, pp. 438–449.

[BKL] L. BABAI, W.M. KANTOR, E.M. LUKS: Computational complexity and the classification of finite simple groups. In: *Proc. 24th FOCS*, IEEE Computer Soc. Press, 1983, pp. 162–171.

[BL] L. BABAI, E.M. LUKS: Canonical labeling of graphs. *15th STOC*, ACM 1983, pp. 171–183.

[Cam] P. J. CAMERON: Finite permutation groups and finite simple groups. *Bull. London Math Soc.* **13** (1981), 1–22.

[FHL] M. L. FURST, J. HOPCROFT, E. M. LUKS: Polynomial-time algorithms for permutation groups. In: *Proc. 21st FOCS*, IEEE Comp. Soc. Press 1980, pp. 36-41.

[FSS] M. FURER, W. SCHNYDER, E. SPECKER: Normal forms for trivalent graphs and graphs of bounded valence. In: Proc. 15th STOC, ACM Press 1983, pp. 161-170.

[Kn] D. E. KNUTH: Efficient representation of perm groups. *Combinatorica* **11** (1991), pp. 57–68.

[Lu1] E.M. LUKS: Isomorphism of graphs of bounded valence can be tested in polynomial time. *J. Comp. Sys. Sci.* **25** (1982), 42–65.

[Lu2] E.M. LUKS: Computing the composition factors of a permutation group in polynomial time. *Combinatorica* **7** (1987), 87–99.

[Lu3] E.M. LUKS: Hypergraph Isomorphism and Structural Equivalence of Boolean Functions. In: *31st STOC*, ACM Press 1999, pp. 652–658.

[Py] L. PYBER: The orders of doubly transitive groups, elementary estimates. *J. Comb. Theory, Ser. A* **62** (1993), 361–366.

[Se] Á. SERESS: *Permutation Group Algorithms.* Cambridge U. Press, 2003.

[Si1] C.C. SIMS: Computation with Permutation Groups. In: *Proc. 2^{nd} Symp. Symb. Algeb. Manip.* ACM, New York, 1971, pp. 23–28.

[Si2] C.C. SIMS: Some group theoretic algorithms. In: *Lecture Notes in Math.* Vol. 697, Springer, 1978, pp. 108-124.

[Sp] D. A. SPIELMAN: Faster Isomorphism Testing of Strongly regular Graphs. In: *Proc. 28th STOC*, ACM Press 1996, pp. 576–584.

[Wi] H. WIELANDT: *Finite Permutation Groups.* Acad. Press, New York 1964.