McMaster University

Advanced Optimization Laboratory

Title:

Bounded-Support Additive Latin Transversals

Authors:

Antoine Deza, Yan Gerard, Yijun Ma, and Sebastian Pokutta

AdvOL-Report No. 2026/4

July 2026, Hamilton, Ontario, Canada

Bounded-Support Additive Latin Transversals

Antoine Deza
McMaster University, Hamilton, Ontario, Canada
deza@mcmaster.ca

Yan Gerard
Université Clermont Auvergne, LIMOS, France
yan.gerard@uca.fr

Yijun Ma
McMaster University, Hamilton, Ontario, Canada
yijun@mcmaster.ca

Sebastian Pokutta
Zuse Institute Berlin and Technische Universität Berlin
pokutta@zib.de

August 7, 2026

Abstract

Let $A = (a_1, \ldots, a_k) \in (\mathbb{Z}_m)^k$, let $B \subseteq \mathbb{Z}_m$ have cardinality k, and write $s = |\operatorname{supp}(A)|$. The additive-transversal problem seeks an ordering $b_1, \ldots, b_k$ of B such that the sums $a_i + b_i$ are pairwise distinct. We reduce this problem to exact-weight perfect matching. The support values of A define edge colors in a bipartite graph, and the prescribed multiplicities of all but one color are encoded as a single target weight in base $k+1$. Applying the exact-weight matching algorithm of Lassota, Łukasiewicz, and Polak [11, Theorem 2] yields a randomized algorithm with running time $(k+1)^{s-1} \operatorname{poly}(k, s, \log m)$. Thus the problem is solvable in randomized polynomial time for every fixed support size s. The reduction applies over arbitrary cyclic groups and does not use the primality assumption appearing in the motivating existence theorem of Alon.

1 Introduction

In a January 2022 blog post, Eppstein highlighted the following open algorithmic question from a talk of Alon [4]. Fix a prime p, a sequence $A = (a_1, \ldots, a_k) \in (\mathbb{Z}_p)^k$ with $k < p$, and a k-subset $B \subseteq \mathbb{Z}_p$. Can one order the elements of B as $b_1, \ldots, b_k$ so that the sums $a_i + b_i$ are pairwise distinct? Alon proved that such an ordering always exists [1], but no polynomial-time construction is known. Consider the general case of an arbitrary cyclic group $\mathbb{Z}_m$. Over $\mathbb{R}$, pairwise distinct sums can be obtained by ordering the elements of A and B increasingly. The full-size case where $k = m$ is governed by Hall's theorem [8][1]: For a sequence $A = (a_1, \ldots, a_m) \in (\mathbb{Z}_m)^m$, there is an ordering $b_1, \ldots, b_m$ of all elements of $\mathbb{Z}_m$ such that the sums $a_i + b_i$ are pairwise distinct if and only if $\sum_{i=1}^m a_i \equiv 0 \pmod m$. Hall's theorem extends to the case when A and B have size $k = m-1$. Hence the cases $k = m$ and $k = m-1$ are completely characterized with the important remark that the

[1]Hall's theorem appears in the mathematics of juggling: in the language of periodic juggling patterns, it provides a necessary and sufficient condition for the feasibility of a prescribed throw sequence.

solutions are constructible in polynomial time [8]. Apart from the brute-force polynomiality for fixed k, the cases $k = m$ and $k = m - 1$ are the only unrestricted regimes for which a polynomial-time construction is known.

Let $\mathrm{supp}(A) = \{\alpha_0, \ldots, \alpha_{s-1}\}$ be the set of distinct values appearing in A. Our principal contribution is a direct reduction of bounded-support additive transversals to exact-weight perfect matching. Noticing that all but one of the support multiplicities can be encoded as digits of a single target weight in base $k + 1$, and applying the exact-weight matching algorithm of Lassota, Łukasiewicz, and Polak [11, Theorem 2] gives randomized search time $(k + 1)^{s-1} \mathrm{poly}(k, s, \log m)$. Thus, additive transversals are randomized polynomial-time constructible for every fixed support size. We use *Color-Counted Matching* as an interface for the reduction: given an edge-colored graph and target counts $r_0, \ldots, r_{h-1}$, find a matching containing exactly r_i edges of color i. The additive-transversal problem has a small, direct matching representation to which the known exact-weight algorithm applies. A more self-contained route through Exact Red Matching and the algorithm of Mulmuley, Vazirani, and Vazirani [12] is also presented.

Organization. Section 2 recalls the relevant background on additive transversals and exact matching. Section 3 states the main results on additive transversals and Color-Counted Matching, together with several applications. Section 4 proves Theorem 3.2. Finally, Section 5 establishes Theorem 3.6, a self-contained alternative to the exact-weight Color-Counted Matching bound.

2 Related Results

Latin squares and their transversals have a long history in combinatorics. In the additive setting over a cyclic group, Hall's theorem [8] gives a complete characterization in the full-size case. Namely, for a sequence $A = (a_1, \ldots, a_m) \in (\mathbb{Z}_m)^m$, there exists an ordering $b_1, \ldots, b_m$ of all elements of $\mathbb{Z}_m$ such that the sums $a_i + b_i$ are pairwise distinct if and only if $\sum_{i=1}^{m} a_i \equiv 0 \pmod{m}$. Moreover, Hall's theorem yields a polynomial-time construction. The near-full case $k = m - 1$ also follows from Hall's theorem and is likewise polynomial-time constructible.

Alon [1] proved in 2000 that, for prime p, every sequence $a_1, \ldots, a_k \in \mathbb{Z}_p$ with $k < p$ and every k-subset $B \subseteq \mathbb{Z}_p$ admits an ordering $b_1, \ldots, b_k$ such that the sums $a_i + b_i$ are pairwise distinct. His proof uses the Combinatorial Nullstellensatz applied to the polynomial

$$\prod_{1 \le i < j \le k} (x_i - x_j)(a_i + x_i - a_j - x_j).$$

Dasgupta, Károlyi, Serra, and Szegedy extended this existence phenomenon to $\mathbb{Z}_{p^\alpha}$ and to $(\mathbb{Z}_p)^\alpha$ [2]. Gao and Wang later gave an alternative proof using group rings [5]. These results provide a strong existence theory, but they do not give a polynomial-time construction for the general repeated-sequence problem. Apart from the brute-force polynomiality when $k = O(1)$ and the full or near-full cases $k = m$ and $k = m - 1$, no general polynomial-time construction is known.

Alon's and Hall's additive-transversal theorems are naturally connected with discrete tomography. In discrete tomography one studies inverse problems in which a finite set of grid points, or equivalently a binary matrix, is reconstructed from prescribed line sums in several directions; these line sums are usually called X-rays [6, 7, 9].

From this viewpoint, Hall's theorem on abelian groups [8], Del Lungo's reconstruction theorem for sums of permutation matrices from diagonal sums [3], and Alon's theorem can all be interpreted as periodic X-ray consistency results for permutation-type incidence structures. The chosen pairs form a periodic binary incidence matrix: one projection prescribes the selected elements of B, another projection prescribes the multiset A, and the all-different condition asks that the third projection be binary. This geometric viewpoint, which links our reconstruction problem to discrete tomography, is illustrated in Fig. 1. In this illustration, the multiset is $A = (0, 2, 2, 3, 5)$ and the set is $B = \{0, 1, 4, 5, 6\}$ in $\mathbb{Z}_7$. Rows are indexed by elements of $\mathbb{Z}_7$ and encode the choice of elements of B. A row indexed by $b \in B$ must contain one selected cell, and a row indexed by $b \notin B$ must contain none. Columns encode the multiplicities in A. The column indexed by a must contain as many selected cells as the multiplicity of a in A. Thus, in the example, the column indexed by 0 contains one selected cell, since 0 has multiplicity 1 in A. The column indexed by 1 contains no selected cell, since 1 has multiplicity 0 in A and the column indexed by 2 contains two selected cells, since 2 has multiplicity 2 in A. Finally, cyclic diagonals encode sums modulo 7. Two selected cells lie on the same diagonal precisely when the corresponding sums $a+b$ are equal. Hence the additive-transversal condition is exactly the requirement that the selected cells occupy distinct diagonals.

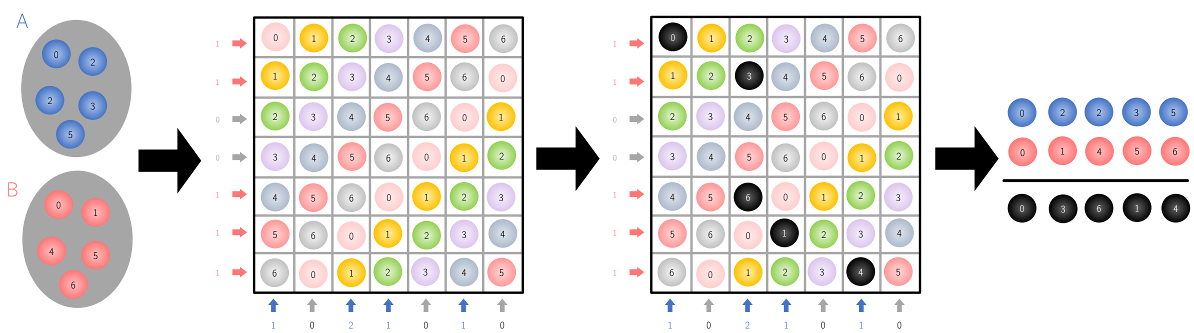

Figure 1: An additive Latin transversal instance viewed as a three-direction discrete tomography instance.

This perspective explains why matching enters the picture. Thus the bounded-support additive-transversal problem may be viewed as a particularly structured instance of a broader discrete-tomography theme: reconstruction from three projections when one projection has bounded support.

On the algorithmic side, Color-Counted Matching is a primitive that appears naturally in the additive-transversal problem but is not specific to it. It includes Exact Matching as the two-color case, and it captures several exact-composition variants of matching problems. For instance, if all target counts are 1, it includes exact rainbow matching. There is also a direct Color-Counted matching formulation of 3-dimensional matching: a triple (x, y, z) may be represented as an edge xy colored by z, and requiring prescribed counts for the z-colors enforces the corresponding constraints in the third coordinate. More generally, Color-Counted Matching captures colored perfect matchings arising in discrete tomography, tilings, permutation-matrix reconstruction, and cycle-cover problems. Several applications of this formulation are described below.

The algorithmic bridge used in this paper is the reduction from Color-Counted Matching to Exact-Weight Perfect Matching. The latter problem asks, given a non-negatively weighted multigraph (V, E) and a target value t, whether (V, E) has a perfect matching whose weight sum is t. Lassota, Łukasiewicz, and Polak [11] give a randomized Monte Carlo algorithm for this problem running in time $\widetilde{O}(t|V|^8 + |E|)$. This is the exact-weight matching engine used in our main reduction. The broader matching-ILP framework of Lassota and Ligthart [10] applies to matching problems with additional linear constraints and therefore also contains prescribed color counts as a special case. Their framework gives other randomized XP algorithms parameterized by the number of additional constraints.

3 Main Results

We now state the two algorithmic problems and the corresponding complexity bounds.

Definition 3.1. Fix integers $m \geq 1$ and $1 \leq k \leq m$. An Additive Transversal instance consists of a sequence $A = (a_1, \dots, a_k) \in (\mathbb{Z}_m)^k$ and a set $B \subseteq \mathbb{Z}_m$ of cardinality k. A *solution* is an ordering $b_1, \dots, b_k$ of the elements of B such that the sums $a_1 + b_1, \dots, a_k + b_k$ are pairwise distinct in $\mathbb{Z}_m$.

Let $\alpha_0, \dots, \alpha_{s-1}$ be the distinct values appearing in A, and let m_j be the multiplicity of α_j in A. Thus $\sum_{j=0}^{s-1} m_j = k$. Our main result is a randomized algorithm whose running time is polynomial for every fixed support size s, with the support-dependent factor made explicit.

Theorem 3.2. *There is a randomized algorithm such that, on input an additive transversal instance over $\mathbb{Z}_m$ with support size $s = |\operatorname{supp}(A)|$, the algorithm either reports failure or outputs a correct solution. If a solution exists, then the algorithm outputs a solution with probability at least $2/3$ in time $(k+1)^{s-1} \operatorname{poly}(k, s, \log m)$. The failure probability can be reduced by independent repetition. In particular, the problem is randomized polynomial-time solvable for every fixed support size s.*

The proof of Theorem 3.2, given in Section 4, proceeds through a more general matching primitive, which we call Color-Counted Matching.

Definition 3.3. A Color-Counted Matching instance consists of a graph (V, E), an edge coloring $\chi : E \to \{0, \dots, h-1\}$, and nonnegative integer target counts $r_0, \dots, r_{h-1}$. The task is to find a matching M such that $|M \cap \chi^{-1}(i)| = r_i$ for every $i \in \{0, \dots, h-1\}$.

We denote $q = \sum_{i=0}^{h-1} r_i$. Every feasible solution of a Color-Counted Matching instance has exactly q edges. Thus, if $|V| = 2q$, every feasible solution is automatically a perfect matching. Similarly, in a bipartite graph $(L \sqcup R, E)$, if $|L| = q$, then every feasible solution automatically saturates L, while vertices of R may remain unmatched.

We now state the complexity of the Color-Counted Matching algorithm obtained from the exact-weight Monte Carlo algorithm of Lassota, Łukasiewicz, and Polak [11]. In particular, this gives randomized polynomial time for every fixed number of colors.

Theorem 3.4. *For $q \geq 1$ and $h \geq 2$, feasibility of a Color-Counted Matching instance can be decided with one-sided error in time $\widetilde{O}\big(q(q+1)^{h-2}|V|^8 + |V|^2 + |E|\big)$. For every $h \geq 1$, a matching with the prescribed counts, whenever one exists, can be constructed with success probability at least $2/3$ in time $(q+1)^{h-1} \operatorname{poly}(|V| + |E|)$.*

Proof. The case $q = 0$ is solved by the empty matching, and the case $h = 1$ reduces to finding a matching of size q. Assume therefore that $q \geq 1$ and $h \geq 2$. If $2q > |V|$, the instance is infeasible. Otherwise add $|V| - 2q$ dummy vertices, join each dummy vertex to every original vertex by a zero-weight edge, and retain the original edges. Every perfect matching of the augmented graph uses exactly q original edges, and conversely every q-edge matching of the original graph extends to such a perfect matching. The augmented graph has at most $2|V|$ vertices and at most $|E| + |V|^2$ edges. Set $Q = q + 1$. Give every original edge of color $i < h - 1$ weight Q^i and every edge of color $h - 1$ weight zero, and set

$$T = \sum_{i=0}^{h-2} r_i Q^i.$$

Equality with T determines the first $h - 1$ color counts by uniqueness of base-Q digits. The omitted count is then determined by the fact that exactly q original edges are selected. Moreover, $T \leq qQ^{h-2}$.

Applying the weighted exact-matching algorithm of Lassota, Łukasiewicz, and Polak [11, Theorem 2] gives the stated decision bound. Amplifying its one-sided error over the polynomially many calls in the standard edge-fixing self-reduction constructs a witness within the stated search bound; any returned matching is checked deterministically. □

Remark 3.5. Partition a set of q labeled slots into sets S_i with $|S_i| = r_i$. Replace every edge of color i by one parallel copy for each slot $j \in S_i$, and assign that copy the weight $2^q + 2^{j-1}$. The partition encoding of Lassota, Łukasiewicz, and Polak [11, Lemma 9] shows that an exact target of $q2^q + (2^q - 1)$ forces every slot to be used exactly once. Together with their Theorem 2 and the same dummy completion, this gives an alternative $2^q \operatorname{poly}(|V| + |E|)$ randomized search algorithm. Hence one may use the better of $(q+1)^{h-1} \operatorname{poly}(|V| + |E|)$ and $2^q \operatorname{poly}(|V| + |E|)$. For the fixed-support application, where $q = k$ can be large and $h = s$ is fixed, the first bound is the relevant one.

For completeness, we also give a more self-contained route through Exact Red Matching and the algorithm of Mulmuley–Vazirani–Vazirani. This route is conceptually elementary, but it expands weights into unary red-black path gadgets and therefore gives a weaker parameter dependence than the Monte Carlo algorithm of Lassota–Łukasiewicz–Polak.

Theorem 3.6. *Let (V, E) be a graph whose edges are colored with colors $0, \dots, h-1$. Let $r_0, \dots, r_{h-1}$ be prescribed target counts, and set $q = \sum_{i=0}^{h-1} r_i$. There is a randomized algorithm which either reports failure or outputs a matching M such that $|M \cap \chi^{-1}(i)| = r_i$ for every $i \in \{0, \dots, h-1\}$. If such a matching exists, then the algorithm outputs one with probability at least $2/3$ in time $\big(|V|^2 + |E|(q+1)^{h-1}\big)^{O(1)}$. The failure probability can be reduced by independent repetition. In particular, since any feasible matching has size q, the problem is randomized polynomial-time solvable for every fixed number of colors h. If the input graph is simple, then $|E| \le |V|^2$ and $q \le \frac{|V|}{2}$ on feasible instances, so the running time is bounded by $(|V| + |E|)^{O(h)}$.*

The algorithms in Theorems 3.2, 3.4, and 3.6 are randomized because they rely on randomized exact-matching primitives. Theorem 3.4 uses the exact-weight matching algorithm of Lassota, Łukasiewicz, and Polak, while the self-contained route leading to Theorem 3.6 uses the algorithm of Mulmuley–Vazirani–Vazirani, as explained in Section 5. Derandomizing these primitives in the generality needed here is a long-standing open problem.

Beyond additive transversals, Theorems 3.4 and 3.6 give immediate fixed-color algorithms for several problems that can be formulated as Color-Counted Matching instances:

First, the formulation applies to the problem of tiling a region R of the triangular grid with prescribed numbers of lozenges in each orientation (the problem is illustrated in Fig. 2). Lozenge tilings of finite regions of the triangular lattice are equivalent to perfect matchings in the dual honeycomb graph. The three possible lozenge orientations induce a 3-coloring of the edges. Every tiling corresponds to a perfect matching of the dual graph with three colors. Either Theorem 3.4 or Theorem 3.6 therefore gives a randomized algorithm with running time $|R|^{O(1)}$, and hence polynomial time in the number of triangles of R.

Second, directed cycle covers with prescribed arc-type counts reduce directly to Color-Counted Matching. Indeed, a directed cycle cover in a digraph $D = (V, A)$ is equivalent to a perfect matching in the bipartite graph with left and right copies of V, where an arc $u \to v$ becomes the edge $u_L v_R$. If the arcs have h types, then the resulting Color-Counted Matching instance has $2|V|$ vertices, $|A|$ edges, and target matching size $|V|$. Thus either Theorem 3.4 or Theorem 3.6 gives a randomized algorithm with running time $(|V|^2 + |A||V|^{h-1})^{O(1)}$.

Third, Theorems 3.4 and 3.6 apply to permutation matrices with prescribed diagonal sums. A permutation matrix of order N is a perfect matching in the bipartite graph between rows and

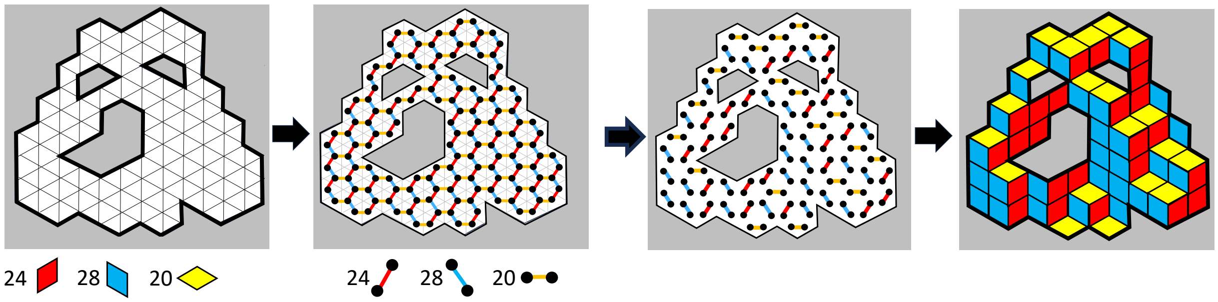

Figure 2: Reduction of the problem of tiling a region of the triangular grid with prescribed numbers of lozenges in each orientation to a Color-Counted Matching problem with three colors.

columns. Coloring an edge (i, j) by its diagonal class turns diagonal sums into color counts. If only h diagonal classes have nonzero prescribed sums, the other diagonals can be deleted, and the resulting bipartite graph has $2N$ vertices, at most Nh edges, and target matching size N. Either Theorem 3.4 or Theorem 3.6 therefore gives a randomized algorithm with running time $N^{O(h)}$, and hence polynomial time for fixed h.

We now turn to the proof of the main additive-transversal theorem.

4 Proof of Theorem 3.2

We now prove Theorem 3.2 by reducing Additive Transversal to Color-Counted Matching and then applying Theorem 3.4. The self-contained Theorem 3.6 gives an alternative route to randomized polynomial time for fixed support size, but the sharper bound stated in Theorem 3.2 comes from Theorem 3.4.

The reduction works as follows. We consider an instance with a multiset $A = (a_1, \ldots, a_k)$ of elements of $\mathbb{Z}_m$ and a set $B \subseteq \mathbb{Z}_m$ of cardinality k. Let $\operatorname{supp}(A) = \{\alpha_0, \ldots, \alpha_{s-1}\}$, and let m_j be the multiplicity of α_j in A.

We build a bipartite graph, called the *diagonal graph* of A and B, as follows. The left vertex set is a copy of B, denoted by L. The right vertex set is a disjoint copy R of the set $\{\alpha_j + b : 0 \leq j < s,\ b \in B\} \subseteq \mathbb{Z}_m$. For every $b \in B$ and every $j \in \{0, \ldots, s-1\}$, we add the edge connecting the left copy of b to the right copy of $b + \alpha_j$, and color it by j. Thus the diagonal graph has k left vertices, at most ks right vertices, and exactly ks edges.

Lemma 4.1. *The additive transversal instance with multiset A and set B has a solution if and only if the diagonal graph built from A and B has a matching using exactly m_j edges of color j for every $j \in \{0, \ldots, s-1\}$.*

Proof. We first assume that the additive transversal problem admits a solution. Thus there is an ordering $b_1, \ldots, b_k$ of the elements of B such that the sums $c_i = a_i + b_i$ are pairwise distinct.

For each index i from 1 to k, let $j(i)$ be the unique index such that $a_i = \alpha_{j(i)}$. By construction of the diagonal graph, there is an edge from the left vertex b_i to the right vertex $c_i = b_i + \alpha_{j(i)}$, and this edge has color $j(i)$. Let $M = \{\, b_i c_i : 1 \leq i \leq k \,\}$. Since the vertices $b_1, \ldots, b_k$ are precisely the elements of B, the matching M saturates L. Moreover, since the sums $c_1, \ldots, c_k$ are pairwise distinct, no two edges of M have the same right endpoint. Hence M is indeed a matching.

It remains to check the color counts. For each color j, the number of edges of color j in M is exactly the number of indices i such that $a_i = \alpha_j$. This number is precisely the multiplicity m_j of α_j in A. Therefore M is a matching using exactly m_j edges of color j for every $j \in \{0, \dots, s-1\}$.

Conversely, assume that the diagonal graph admits a matching M using exactly m_j edges of color j for every $j \in \{0, \dots, s-1\}$. Since $\sum_{j=0}^{s-1} m_j = k$, the matching M has exactly k edges, and therefore saturates the left side L.

For each $b \in B$, let c_b be the right endpoint matched to the left copy of b. If the edge bc_b has color j, then by construction $c_b = b + \alpha_j$.

For each color j, let $I_j = \{i \in \{1, \dots, k\} : a_i = \alpha_j\}$ and let B_j be the set of elements $b \in B$ whose left copy is matched by an edge of color j. By the prescribed color counts, both I_j and B_j have size m_j. Choose an arbitrary bijection from I_j to B_j. Doing this for every j gives a bijection between the index set $\{1, \dots, k\}$ of A and the set B. Equivalently, it gives an ordering $b_1, \dots, b_k$ of B.

If $i \in I_j$, then $a_i = \alpha_j$, and the matched right endpoint of the left copy of b_i is $b_i + \alpha_j = b_i + a_i$. Since M is a matching, these right endpoints are pairwise distinct. Therefore the sums $a_i + b_i$ are pairwise distinct, and the additive transversal instance has a solution. □

We now use Lemma 4.1 and the Color-Counted Matching algorithm to prove Theorem 3.2.

Proof. We consider an additive transversal instance over $\mathbb{Z}_m$, with $|A| = |B| = k$, and we denote $s = |\operatorname{supp}(A)|$. By Lemma 4.1, the instance reduces to a Color-Counted Matching instance on the diagonal graph, with target counts $m_0, \dots, m_{s-1}$. Hence the target matching size is $q = \sum_{j=0}^{s-1} m_j = k$, and the number of colors is $h = s$. The diagonal graph has k left vertices, at most ks right vertices, and exactly ks edges. Applying Theorem 3.4 with $q = k$ and $h = s$ gives randomized search time $(k+1)^{s-1} \operatorname{poly}(k, s)$ in the unit-cost model.

It remains to account for arithmetic in $\mathbb{Z}_m$. Constructing the diagonal graph requires computing, comparing, and storing the ks residues $b + \alpha_j \pmod m$, each represented with $O(\log m)$ bits. The matching weights use $O(s \log(k+1))$ bits. Consequently, the complete bit complexity is $(k+1)^{s-1} \operatorname{poly}(k, s, \log m)$.

If the Color-Counted Matching algorithm returns a matching with the prescribed color-count vector, Lemma 4.1 converts it into a valid additive transversal. The returned solution can be checked deterministically. If a solution exists, the success probability is at least $2/3$, and it can be amplified by independent repetition. This proves Theorem 3.2. □

5 Proof of Theorem 3.6

We now provide the more self-contained proof leading to Theorem 3.6.

5.1 From h colors to two colors

We first reduce Color-Counted Matching to Exact Red Matching.

Definition 5.1. An Exact Red Matching instance consists of a graph (V, E), a set of red edges $E_{\text{red}} \subseteq E$, and an integer T. The task is to find a perfect matching M such that $|M \cap E_{\text{red}}| = T$.

The Exact Red Matching problem is the special case of Color-Counted Matching with two colors, red and black, in which the matching is required to be perfect. Indeed, every perfect matching has $\frac{|V|}{2}$ edges. Thus, prescribing the number of red edges of a perfect matching is equivalent to prescribing target counts r_{red} and r_{black} satisfying $r_{\text{red}} + r_{\text{black}} = \frac{|V|}{2}$.

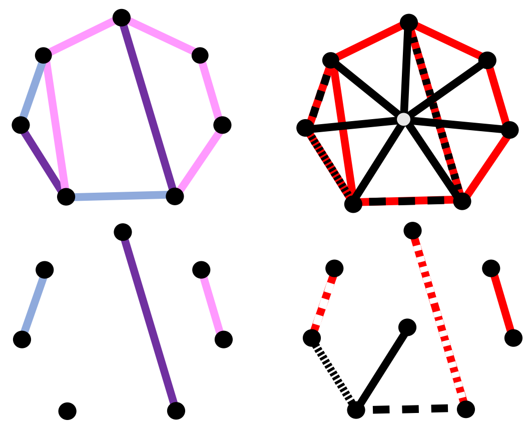

Figure 3: **Reduction of a 3-color Color-Counted Matching instance with target counts $(1, 1, 1)$ to an Exact Red Matching instance with 21 red edges.** The Color-Counted Matching instance appears on the left, and the red-black Q-augmented graph appears on the right. The target matching size is $q = 1 + 1 + 1 = 3$, so $Q = q + 1 = 4$, and the red target is $T = Q^0 + Q^1 + Q^2 = 1 + 4 + 16 = 21$. Corresponding solutions are shown below.

The red-black Q-augmented graph. We now describe the construction of the Exact Red Matching instance. Let $H = (V, E)$ be a Color-Counted Matching instance whose edges are colored with colors $0, \ldots, h-1$, and let $r_0, \ldots, r_{h-1}$ be the prescribed target counts. We denote $q = \sum_{i=0}^{h-1} r_i$. Thus any matching satisfying the prescribed color counts has size q.

If $q > \lfloor \frac{|V|}{2} \rfloor$, then no matching of size q exists in an $|V|$-vertex graph, and we may immediately report failure. Hence we assume from now on that $q \leq \lfloor \frac{|V|}{2} \rfloor$. If $q = 0$, the empty matching is a solution of the Color-Counted Matching instance. We therefore assume that $q \geq 1$ and set $Q = q + 1$.

The red-black Q-augmented graph $H' = (V', E')$ is built from H as follows. The construction is illustrated in Fig. 3.

1. We initialize $V' = V$ and $E' = \emptyset$.
2. Each edge $e = (u, v)$ of color i in E is replaced by a path $u = x_{e,0},\ x_{e,1},\ \ldots,\ x_{e,2Q^i-1} = v$ with $2Q^i - 1$ edges and $2Q^i - 2$ new internal vertices added to V' and used only for the edge e.

   The Q^i edges $x_{e,2t}x_{e,2t+1}$ with $0 \leq t < Q^i$ are colored red, while the $Q^i - 1$ edges $x_{e,2t+1}x_{e,2t+2}$ with $0 \leq t < Q^i - 1$ are colored black. All these path edges are added to E'.
3. If $2q < |V|$, then we add to V' a set D of $|V| - 2q$ dummy vertices and join every vertex of D to every original vertex of V by a black edge. These dummy edges are added to E'.

The target number of red edges in the Exact Red Matching instance is $T = \sum_{i=0}^{h-1} r_i Q^i$.

Size of the red-black Q-augmented graph. We now compute the size of the red-black Q-augmented graph. For an edge $e \in E$ of color $\chi(e)$, the replacement path has $2Q^{\chi(e)}$ vertices and $2Q^{\chi(e)} - 1$ edges. Since the two endpoints of the path are the original endpoints of the edge, this replacement adds $2Q^{\chi(e)} - 2$ new vertices. Therefore $|V'| = |V| + (|V| - 2q) + \sum_{e \in E}(2Q^{\chi(e)} - 2)$, and hence $|V'| \leq 2|V| + 2|E|Q^{h-1}$.

Similarly, the path replacements contribute $\sum_{e \in E}(2Q^{\chi(e)} - 1)$ edges. The dummy vertices contribute $|V|(|V| - 2q)$ black edges, since each dummy vertex is joined to every original vertex. Hence $|E'| = |V|(|V| - 2q) + \sum_{e \in E}(2Q^{\chi(e)} - 1)$, and therefore $|E'| \leq |V|^2 + 2|E|Q^{h-1}$.

Thus the red-black Q-augmented graph has size $|H'| = O\big(|V|^2 + |E|Q^{h-1}\big)$. Since $Q = q+1$, this is $|H'| = O\big(|V|^2 + |E|(q+1)^{h-1}\big)$.

Remark 5.2. The construction of the red-black Q-augmented graph is not necessarily optimal. One could use mixed radices instead of the uniform base Q. More precisely, if Q_i is a strict upper bound on the number of selected edges of color i, one may define $W_0 = 1$ and $W_i = \prod_{j=0}^{i-1} Q_j$ for $i \geq 1$, and replace the weight Q^i by W_i. A preprocessing step could further reduce the size of the augmented graph by taking the target counts r_i into account. We do not use this optimization here.

Equivalence between h colors and two colors in the Q-augmented graph. We now show the equivalence between the Color-Counted Matching instance in H and the Exact Red Matching instance in the red-black Q-augmented graph H'.

Proposition 5.3. *Let $H = (V, E)$ be a Color-Counted Matching instance with colors $0, \dots, h-1$ and target color counts $r_0, \dots, r_{h-1}$. Let $q = \sum_{i=0}^{h-1} r_i$, assume that $1 \leq q \leq \lfloor \frac{|V|}{2} \rfloor$, and set $Q = q+1$. Let $H' = (V', E')$ be the red-black Q-augmented graph constructed above. Then H has a matching with the prescribed color counts if and only if H' has a perfect matching with exactly $T = \sum_{i=0}^{h-1} r_i Q^i$ red edges.*

Proof. We first prove that, given a solution M of the Color-Counted Matching instance, we can build a solution of the Exact Red Matching instance for the corresponding red-black Q-augmented graph with exactly $T = \sum_{i=0}^{h-1} r_i Q^i$ red edges. Thus M contains exactly r_i edges of color i for every $i \in \{0, \dots, h-1\}$, and therefore $|M| = q$.

We build a matching M' in H' as follows. For each edge $e = uv$ of M, we select in M' all the red edges of the path replacing e. For each edge $e = uv$ of E that does not belong to M, we select in M' all the black edges of the path replacing e. Finally, we match the $|V| - 2q$ vertices of V that are not matched by M bijectively to the $|V| - 2q$ dummy vertices.

This gives a perfect matching of H'. Indeed, in every path replacing an edge of H, all internal vertices are matched. If the corresponding edge belongs to M, then the two original endpoints are matched through the red edges of the path. If it does not belong to M, then the black edges of the path match all internal vertices and leave the two original endpoints available. These available original vertices are either matched through another selected path or, if they are unmatched in M, matched to dummy vertices.

The red edges selected in M' come only from the paths corresponding to the edges of M. Each selected edge of color i contributes exactly Q^i red edges. Hence the total number of red edges in M' is $\sum_{i=0}^{h-1} r_i Q^i = T$. Thus M' is a perfect matching of H' with exactly T red edges.

We now prove the converse implication. Assume that H' admits a perfect matching M' with exactly T red edges. We build from it a matching M in H.

For each edge $e = uv$ of H, consider the alternating red-black path replacing e in H'. Since the internal vertices of this path have no incident edges outside the path, a perfect matching of H' has only two possible states on this path: either it contains all the red edges of the path, or it contains all the black edges of the path. We call the path active in the first case and inactive in the second case. We select the edge e in M if and only if the path replacing e is active.

It follows that M is a matching in H. Indeed, two active paths cannot share an original endpoint, since otherwise M' would contain two edges incident with the same vertex.

We now observe that M has exactly q edges. In the graph H', every dummy vertex is matched to an original vertex. Since there are $|V| - 2q$ dummy vertices, exactly $|V| - 2q$ original vertices are matched to dummy vertices. The remaining $2q$ original vertices must be matched through active paths. Each active path uses exactly two original endpoints. Hence exactly q paths are active, and therefore $|M| = q$.

For each $i \in \{0, \ldots, h-1\}$, let y_i be the number of selected edges of color i in M. Equivalently, y_i is the number of active paths of weight Q^i in H'. Since an active path of color i contributes exactly Q^i red edges, the number of red edges in M' is $\sum_{i=0}^{h-1} y_i Q^i$. Since M' has exactly T red edges, $\sum_{i=0}^{h-1} y_i Q^i = \sum_{i=0}^{h-1} r_i Q^i$.

We now use uniqueness of base-Q expansion. Since $|M| = q$ and $Q = q + 1$, we have $y_i < Q$ for every i. Also, since $r_i \le q$ and $Q = q + 1$, we have $r_i < Q$ for every i. Thus both sides of the previous equality are base-Q expansions with digits strictly smaller than Q. Consequently, for every $i \in \{0, \ldots, h-1\}$, we have $y_i = r_i$.

Therefore M contains exactly r_i edges of color i, and so M is a solution of the original Color-Counted Matching instance. □

5.2 Proof of Theorem 3.6

The original result of Mulmuley–Vazirani–Vazirani places exact matching in randomized NC. In this paper we only use the resulting randomized polynomial-time search algorithm that we restate here.

Theorem 5.4 ([12]). *Let (V, E) be a graph, let $E_{\text{red}} \subseteq E$, and let $T \ge 0$. There is a randomized algorithm running in time polynomial in $|V| + |E|$ which either reports failure or outputs a perfect matching M such that $|M \cap E_{\text{red}}| = T$. If such a perfect matching exists, then the algorithm outputs one with probability at least $2/3$. If no such perfect matching exists, then the algorithm reports failure with probability 1. The success probability on yes-instances can be increased to $1 - \varepsilon$, for any $\varepsilon > 0$, by independent repetition, with an additional multiplicative factor $O(\log(1/\varepsilon))$ in the running time.*

The original theorem of Mulmuley–Vazirani–Vazirani applies to general graphs. Applying Theorem 5.4 to the Q-augmented red-black graph H' gives a randomized algorithm whose running time is polynomial in $|V'| + |E'|$. We now prove Theorem 3.6.

Proof of Theorem 3.6. Let $q = \sum_{i=0}^{h-1} r_i$. If $q = 0$, the algorithm returns the empty matching. We therefore assume that $q \ge 1$. If $q > \lfloor \frac{|V|}{2} \rfloor$, then no matching of size q exists in (V, E), and the algorithm reports failure. We therefore assume $q \le \lfloor \frac{|V|}{2} \rfloor$, and set $Q = q + 1$.

We construct the red-black Q-augmented graph $H' = (V', E')$ as in Proposition 5.3, and set the red target to be $T = \sum_{i=0}^{h-1} r_i Q^i$. By Proposition 5.3, the original Color-Counted Matching instance has a solution if and only if H' has a perfect matching with exactly T red edges.

We now apply the exact matching theorem of Mulmuley–Vazirani–Vazirani to the red-black graph H' with red-edge set E'_{red} and target T, where E'_{red} is the set of red edges of E'. If such a perfect matching exists, the MVV algorithm outputs one with probability at least $2/3$; if none exists, it reports failure with probability 1. Any returned perfect matching can be checked deterministically.

The running time for building the red-black Q-augmented graph is linear in its size, namely $|V'| + |E'| = O\big(|V|^2 + |E| Q^{h-1}\big)$. The MVV algorithm runs in time polynomial in the size of its input, so the total running time is $\big(|V|^2 + |E| Q^{h-1}\big)^{O(1)}$. Since $Q = q + 1$, this is $\big(|V|^2 + |E|(q+1)^{h-1}\big)^{O(1)}$.

Finally, if the MVV algorithm returns a perfect matching of H' with exactly T red edges, the constructive correspondence in Proposition 5.3 allows us to extract the active path gadgets from it. These active gadgets form a matching of H with color-count vector $(r_0, \ldots, r_{h-1})$. Hence the algorithm outputs a correct Color-Counted Matching solution. The failure probability can be reduced by independent repetition. □

Acknowledgments

We thank Alexandra Lassota for drawing our attention to the ICALP 2022 exact-matching result [11]. Research reported in this paper was partially supported by the Berlin Mathematics Research Center MATH$^+$ (EXC-2046/2, project ID 390685689), funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy as well as by the Natural Sciences and Engineering Research Council of Canada Discovery (Grant program number RGPIN-2026-06553).

References

[1] Noga Alon. Additive Latin transversals. *Israel Journal of Mathematics*, 117(1):125–130, 2000. doi: 10.1007/BF02773567.

[2] Samit Dasgupta, Gyula Károlyi, Oriol Serra, and Balázs Szegedy. Transversals of additive Latin squares. *Israel Journal of Mathematics*, 126(1):17–28, 2001. doi: 10.1007/BF02784149.

[3] Alberto Del Lungo. Reconstructing permutation matrices from diagonal sums. *Theoretical Computer Science*, 281(1–2):235–249, 2002. doi: 10.1016/S0304-3975(02)00015-4.

[4] David Eppstein. Open algorithmic problems from a talk by Alon, January 2022. URL https://11011110.github.io/blog/2022/01/27/open-algorithmic-problems.html. Blog post.

[5] Weidong Gao and D. J. Wang. Additive Latin transversals and group rings. *Israel Journal of Mathematics*, 140(1):375–380, 2004. doi: 10.1007/BF02786641.

[6] Richard J. Gardner and Peter Gritzmann. Discrete tomography: Determination of finite sets by X-rays. *Transactions of the American Mathematical Society*, 349(6):2271–2295, 1997. doi: 10.1090/S0002-9947-97-01741-8.

[7] Richard J. Gardner, Peter Gritzmann, and Doris Prangenberg. On the computational complexity of reconstructing lattice sets from their X-rays. *Discrete Mathematics*, 202(1–3):45–71, 1999. doi: 10.1016/S0012-365X(98)00347-1.

[8] Marshall Hall. A combinatorial problem on abelian groups. *Proceedings of the American Mathematical Society*, 3(4):584–587, 1952. doi: 10.1090/S0002-9939-1952-0050579-7.

[9] Gabor T. Herman and Attila Kuba, editors. *Discrete Tomography: Foundations, Algorithms, and Applications*. Applied and Numerical Harmonic Analysis. Birkhäuser, Boston, MA, 1999. ISBN 978-0-8176-4101-6. doi: 10.1007/978-1-4612-1568-4.

[10] Alexandra Lassota and Koen Ligthart. Parameterized algorithms for matching integer programs with additional rows and columns. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, volume 334 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 112:1–112:18, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.ICALP.2025.112. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2025.112.

[11] Alexandra Lassota, Aleksander Łukasiewicz, and Adam Polak. Tight vector bin packing with few small items via fast exact matching in multigraphs. In Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 87:1–87:15, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.ICALP.2022.87. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2022.87.

[12] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987. doi: 10.1007/BF02579206.