SHARP PROXIMITY FOR GROUP-LABELED MATROID BASES

AVINASH BHARDWAJ

Abstract. Let the elements of a matroid be labeled in an abelian group, and let F be a finite set of forbidden basis-label sums. We give an alternative proof of the sharp unweighted proximity bound established by Lizhong Chen: for every representable matroid, every basis is within $|F|$ exchanges of an F-avoiding basis whenever one exists. This holds over arbitrary representation fields, with no compatibility condition on the label group. The bound is sharp even for spanning trees. The proof combines column operations within label classes, induction through a hyperplane, and the classical unique-representation sumset inequality. For weighted proximity, we prove that a polynomial nonvanishing constraint of degree d admits an optimal feasible basis within d exchanges of every unconstrained optimum. For elementary abelian p-group labels and a representation in characteristic p, this gives the exact worst-case weighted radius; for binary labels it is $\lfloor \log_2(|F|+1) \rfloor$. We also establish a sharp group-dependent bound for finite abelian p-groups and extensions to simultaneous constraints and realizable valuated costs.

1. Introduction

An additional constraint can change an optimal solution substantially, even when the original optimization problem is easy. For matroid bases, proximity asks how much of a basis can be retained when the constraint is imposed. There are two versions of the question: find any nearby feasible basis, or find a nearby basis that is cheapest among all feasible ones. We study both, with bounds independent of the rank and ground-set size.

Let $\mathcal{B}(M)$ denote the bases of a finite matroid M on E. For bases A, B, distance counts exchanges:

$$\operatorname{dist}(A, B) = |B \setminus A| = \tfrac{1}{2}|A \triangle B|.$$

Given a labeling $\psi : E \to \Gamma$ in an abelian group, write $\psi(S) = \sum_{e \in S} \psi(e)$. A basis B is *F-avoiding* if $\psi(B) \notin F$, where $F \subseteq \Gamma$ is finite. Hörsch et al. [8] conjecture that, whenever an F-avoiding basis exists, every basis is within $|F|$ exchanges of one. Chen [3, Theorem 1.1] established this bound for every representable matroid as part of a more general theorem for simultaneous label constraints. We give an alternative proof of its single-constraint case.

Theorem 1.1 (Chen [3, Theorem 1.1], single-constraint case)**.** *Let M be representable over a field K, let Γ be any abelian group, let $F \subseteq \Gamma$ be finite, and let $\psi : E \to \Gamma$. If an F-avoiding basis exists, then for every basis A there is an F-avoiding basis B with*

$$\operatorname{dist}(A, B) \leq |F|.$$

The bound is sharp even for graphic matroids.

The group may be finite or infinite, and its torsion need not match the characteristic of K. In particular, the theorem applies to graphic, cographic, and regular matroids with arbitrary abelian labels. Its key assertion is a label-counting statement: if a basis B is the only basis with its label and I is an independent set disjoint from B, then at least $|I| + 1$ different basis labels occur. To prove this, we modify a representation by elementary column operations between equally labeled elements. These operations preserve the attainable basis labels and allow induction through a coordinate hyperplane. Theorem 1.1 concerns unweighted proximity. The general nonrepresentable case and the weighted conjecture for arbitrary abelian labels are

2020 *Mathematics Subject Classification.* 05B35, 90C27, 05E16.

Key words and phrases. Matroid optimization, proximity, group-labeled bases, polynomial method, augmentation ideal.

outside the scope of our proofs; see the note added at the end of the paper for the resolution communicated by Liu, Oki, and Xu.

For the weighted question, assign real costs $w : E \to \mathbb{R}$ and define $w(B) = \sum_{e \in B} w(e)$. We compare an unconstrained minimum-cost basis with a minimum-cost feasible basis. Our second result uses a polynomial to describe feasibility. If K is a field, let $\mathbf{1}_B \in K^E$ be the incidence vector of B. The cost field $\mathbb{R}$ need not coincide with K.

Theorem 1.2. *Let M be representable over a field K, and let $P \in K[x_e : e \in E]$ have total degree at most d. Suppose that a basis B with $P(\mathbf{1}_B) \neq 0$ exists. For every minimum-cost basis A of M, there is a minimum-cost basis B among those satisfying $P(\mathbf{1}_B) \neq 0$ such that* $\operatorname{dist}(A, B) \leq d$.

This theorem includes arbitrary real costs and every choice of the unconstrained optimum, including ties. The bound is sharp even for a connected multigraph. It concerns polynomial *nonvanishing*; a polynomial equality need not have an equivalent nonvanishing description of the same degree.

For elementary abelian labels, we obtain an exact worst-case weighted bound. For a prime p and $r \geq 0$ write

$$r = a(p-1) + b, \qquad 0 \leq b < p-1, \qquad \mu_p(r) = (b+1)p^a.$$

For $k \geq 0$, define the inverse threshold

$$\lambda_p(k) = \max\{r \in \mathbb{Z}_{\geq 0} : \mu_p(r) \leq k+1\}. \tag{1}$$

The function μ_p is strictly increasing. Explicitly, if $a = \lfloor \log_p(k+1) \rfloor$, then

$$\lambda_p(k) = a(p-1) + \left\lfloor \frac{k+1}{p^a} \right\rfloor - 1. \tag{2}$$

In particular $\lambda_2(k) = \lfloor \log_2(k+1) \rfloor$ and $\lambda_p(k) \leq k$.

Theorem 1.3. *Let K have characteristic p, let M be K-representable, and let $\psi : E \to \mathbb{F}_p^m$. If an F-avoiding basis exists, then for every unconstrained minimum-cost basis A there is a minimum-cost F-avoiding basis B with*

$$\operatorname{dist}(A, B) \leq \lambda_p(|F|) \leq \min\{|F|, m(p-1)\}.$$

For every p, every $m \geq 1$, and every $0 \leq k < p^m$, the first bound is attained by a labeling of a connected multigraph with $|F| = k$. The unconstrained and constrained optima in these examples can both be unique.

Thus the weighted $|F|$ bound improves to a logarithmic bound for each fixed prime p. For $p = 2$, the bound applies to every binary matroid. For every p, it applies to all regular matroids, including graphic and cographic matroids. The group dimension does not occur in the bound, except through the necessary inequality $|F| < p^m$.

The next result concerns arbitrary finite abelian p-groups. Here the bound depends on the group rather than on the forbidden-set size. For a set $T \subseteq \Gamma$ of allowed labels, a basis B is *T-labeled* if $\psi(B) \in T$; no further condition is imposed on T.

Theorem 1.4. *Let K have characteristic p, let M be K-representable, and let*

$$\Gamma = \bigoplus_{j=1}^{m} \mathbb{Z}_{p^{a_j}}, \qquad a_j \geq 1, \qquad \nu(\Gamma) = \sum_{j=1}^{m} (p^{a_j} - 1).$$

For every labeling $\psi : E \to \Gamma$ and every set $T \subseteq \Gamma$ for which a T-labeled basis exists, every unconstrained minimum-cost basis lies within $\nu(\Gamma)$ exchanges of some minimum-cost T-labeled basis. This bound is sharp even for connected multigraphs.

The number $\nu(\Gamma) + 1$ is the Davenport constant of an abelian p-group [11]. Our proof of the proximity bound is self-contained and uses only a nilpotence calculation.

A further extension concerns nonadditive costs. Section 8 replaces additive costs by valuations of maximal minors plus arbitrary linear costs. The same degree bound holds for a valuation-penalized objective; when the valuation is trivial on the prime subfield, it also holds for polynomial feasibility. In characteristic p, the sharp bound $\lambda_p(|F|)$ survives this extension.

1.1. **A common proof framework: isolation and separation.** The basic reduction begins by choosing a feasible basis B nearest to the prescribed basis A, with cost optimality imposed first in the weighted setting. Contract $A \cap B$ and delete the elements outside $A \cup B$. The remaining bases A_0, B_0 are disjoint, and the rank r of this minor is exactly the original distance. The choice of B creates an isolation property. Without costs, B_0 is the only feasible basis of the minor. With additive costs and A an unconstrained optimum, it is the only feasible basis whose complement is also a basis (Lemma 3.1). This distinction determines which structural information the rest of the proof may use.

In the unweighted case, equal-label column operations preserve the isolated label while concentrating each label class outside a hyperplane into at most one column. Every basis of the hyperplane extends by every outside column. A sumset inequality then turns isolation into a lower bound on the number of basis labels.

In the weighted case, the relevant family consists only of complementary bases. A determinant built from a representation and its dual gives a linear combination of their evaluations that vanishes for every polynomial of degree less than r. A polynomial separating the isolated feasible basis would leave exactly one nonzero term, a contradiction. Bounds on separator degree, or equivalently on support in the augmentation filtration, yield the label-specific weighted conclusions. For valuated costs, a unique term of smallest valuation replaces a unique nonzero term.

These arguments share the isolation reduction but use different certificates against isolation. For comparison, Appendix A constructs three compatible exchange families for two spanning trees. Alternating finite differences on such a family give an elementary version of the vanishing-moment identity (Lemma 4.1).

1.2. **Related work and the role of the present results.** Liu and Xu [10] introduced closeness and strong closeness for group-constrained bases. In particular, they prove the weighted exact-label bound $D(\Gamma) - 1$ for strongly base-orderable matroids and finite abelian groups Γ, where $D(\Gamma)$ is the Davenport constant [10, Theorem 6.1, full version]. They also obtain weighted results for small groups. Their strong-block-isolation reduction [10, Proposition 6.2, full version] supplies the weighted principle used here; we state its extension to an arbitrary feasible family.

Hörsch, Imolay, Mizutani, Oki, and Schwarcz [8] formulate the $|F|$ proximity conjecture and its weighted version, prove a cubic unweighted bound for graphic matroids, and give group-ring determinant algorithms. The weighted $|F|$ bound is already known for $|F| = 1$ in arbitrary matroids and for arbitrary finite F in strongly base-orderable matroids [8, Lemma 3.2 and the discussion following Conjecture 7.1, full version]. These forbidden-set bounds should be distinguished from the exact-label Davenport bound above.

Their counterexample to the unrestricted Davenport bound uses $\Gamma = \mathbb{Z}_2^d$, $d \geq 4$, and a rank-$(2^{d-1} - 1)$ sparse paving matroid on 2^d elements [8, Theorem 6.14 and Lemma 6.16, full version]. Remark 5.3 describes the rank-seven instance and exhibits a Vámos minor. It is therefore not representable over any field and does not contradict Theorem 1.4. The general determinant encoding is already present in their work.

Bérczi and Schwarcz [2] give another relaxation of strong base orderability. For two disjoint graphic bases, they construct an alternating path on their elements such that every circuit in their union contains two consecutive elements [2, Theorem 12, preprint version]. Our structural certificate is instead a partition of one prescribed tree into three families, each allowing all subset exchanges with distinct partners in the other tree. We distinguish this certificate from an ordering that only guarantees particular exchange sequences.

Garamvölgyi, Mizutani, Oki, Schwarcz, and Yamaguchi [6] prove the unweighted $|F|$ conjecture for sparse paving matroids and for $|F| \leq 4$, and study several simultaneous label constraints. Their results permit arbitrary abelian labels. Chen [3], in a preprint deposited on 22 September 2026, establishes the full unweighted bound for representable matroids over every field. More generally, for simultaneous forbidden sets F_t in arbitrary abelian label groups, his Theorem 1.1 gives the bound $\prod_t(|F_t|+1)-1$. Thus Theorem 1.1 is already contained in his work. His proof uses exterior contraction and interpolation over group algebras; the proof given here uses label-preserving column operations and hyperplane induction. Our sharper weighted bounds for elementary abelian groups require compatibility between the group and the representation field.

The additive-combinatorial ingredient in the unweighted proof is the classical Scherk–Kemperman inequality: a uniquely represented sum forces $|U+D| \geq |U|+|D|-1$. We use its form for finite subsets of an arbitrary abelian group; see [7]. The proximity-to-isolation reduction is standard in the preceding matroid literature. The additional step here is to normalize a representation within each label class, preserving its basis-label set before applying hyperplane induction.

Eisenbrand, Rohwedder, and Węgrzycki [5] establish sensitivity, proximity, and fixed-parameter algorithms for exact integer-weight matroid problems. Their comparison with a constrained polyhedral solution differs from the comparison with an unconstrained optimal basis made here.

The classical polynomial ingredients also deserve explicit separation from the matroid conclusions. Polynomial separation on a Boolean cube is closely related to the hyperplane-covering theorem of Alon and Füredi [1]. The support estimate used for Theorem 1.3 is the minimum-distance bound for generalized Reed–Muller codes [4]; see also [9]. We give an elementary proof in the form required here.

The algebraic arguments combine these ingredients through a determinant built from a representation and its *dual*. Their orthogonality gives vanishing moments on complementary bases. Over an elementary abelian group, the same determinant belongs to a power of the augmentation ideal; the classical support bound then translates its rank into a lower bound on the number of labels. The weighted reduction completes the passage from those identities to optimal bases.

1.3. **Organization.** Section 2 proves the sharp unweighted bound by label-preserving column operations and hyperplane induction. Section 3 gives the weighted reduction and the representation–dual determinant. Section 4 proves polynomial proximity. Sections 5 and 6 establish the group-algebra and sharp elementary-group results. Section 7 gives consequences, enumeration algorithms, and sharp examples. Section 8 treats realizable valuated costs. Section 9 discusses the limitations of our approach to weighted arbitrary-group proximity, in light of the communication recorded in the closing note. Appendix A gives the separate structural result on compatible exchanges in spanning trees.

2. Sharp unweighted proximity

We first prove the label-counting statement that drives the argument. Its independent set need not be a basis; this stronger formulation is what makes induction through a hyperplane possible.

For a labeled matroid M, write $L_\psi(M) = \{\psi(C) : C \in \mathcal{B}(M)\}$ for its set of basis labels. A basis B is *isolated* if no other basis has label $\psi(B)$.

We use the classical Scherk–Kemperman inequality: if finite nonempty sets $U, D \subseteq \Gamma$ have an element with exactly one representation as $u+d$, with $(u,d) \in U \times D$, then

$$|U+D| \geq |U|+|D|-1. \tag{3}$$

This statement applies to arbitrary abelian groups, including infinite groups; only the two sets must be finite. See, for example, Hamidoune [7] for the more general inequality with the number of representations subtracted in place of 1.

Lemma 2.1. *Let M be represented by vectors $(v_e : e \in E)$ over a field K. For distinct elements a, c with $\psi(a) = \psi(c)$, replace v_c by $v_c - tv_a$, where $t \in K$, leaving all other columns unchanged. The resulting matroid has exactly the same set of basis labels. If a, c lie in an independent set I, that set remains independent. If an isolated basis B is disjoint from $\{a, c\}$, it remains an isolated basis.*

Proof. Use a full-row-rank representation. An invertible column operation preserves the rank of the whole matrix. Consider a basis C of the new matrix. If $c \notin C$, its columns did not change. If $a, c \in C$, its determinant did not change. Otherwise, multilinearity in column c shows that at least one of C and $C - c + a$ was a basis of the old matrix. These two sets have the same label. Thus every new basis label was already attained. Applying the inverse operation gives the reverse inclusion.

The operation restricted to I is invertible, so I stays independent. The columns of B do not change. If $C \neq B$ is a new basis, one of the old bases just exhibited has its label and is different from B: in the only case requiring expansion, both candidate sets contain an element of $\{a, c\}$, which is disjoint from B. Isolation follows. □

Theorem 2.2. *Let M be representable over a field K, let Γ be an abelian group, and let $\psi : E(M) \to \Gamma$. If B is an isolated basis and $I \subseteq E(M) \setminus B$ is independent, then*

$$|L_\psi(M)| \geq |I| + 1.$$

Proof. We induct on $r = \operatorname{rank}(M)$. The assertion is immediate when $r = 0$. We may restrict to $B \cup I$: isolation and independence persist, and this can only remove basis labels. Write $k = |I|$, and use the columns of $B = \{b_1, \ldots, b_r\}$ as coordinates, so that they form the identity matrix.

Within each label class of I, consider the columns whose first coordinate is nonzero. Choose one such column a, if any. For every other such column c in that class perform the operation

$$v_c \longleftarrow v_c - \frac{(v_c)_1}{(v_a)_1} v_a.$$

The chosen column remains unchanged, and the operation makes the first coordinate of c zero. By Lemma 2.1, throughout these operations B remains isolated, I remains independent, and the set of basis labels is unchanged. Denote the resulting matroid by N.

Let $J \subseteq I$ consist of the columns now having nonzero first coordinate, and let $\ell = |J|$. The labels of the elements of J are pairwise distinct. None equals $\psi(b_1)$: otherwise $B - b_1 + j$ would be a second basis with label $\psi(B)$. Consequently

$$D = \{\psi(b_1)\} \cup \{\psi(j) : j \in J\} \quad \text{has size } \ell + 1.$$

The columns in

$$E_0 = (B \setminus \{b_1\}) \cup (I \setminus J)$$

lie in the coordinate hyperplane $x_1 = 0$ and span it. Thus $N_0 = N|E_0$ has rank $r - 1$, $B_0 = B \setminus \{b_1\}$ is a basis, and $I \setminus J$ is an independent set of size $k - \ell$ disjoint from B_0. Every basis of N_0 extends to a basis of N by adding b_1. It follows that B_0 is isolated in N_0. The induction hypothesis, applied over the same field, gives

$$U = L_\psi(N_0), \qquad |U| \geq k - \ell + 1.$$

Every column in $\{b_1\} \cup J$ lies outside the hyperplane, so every basis of N_0 extends to a basis of N by adding any one of them. Therefore $U + D \subseteq L_\psi(N)$. Moreover, $\psi(B)$ has exactly one representation in $U + D$, namely

$$\psi(B) = \psi(B_0) + \psi(b_1).$$

Indeed, any representation using the label of an element of J would produce a basis different from B with the same label; the labels in D are distinct. A representation using $\psi(b_1)$ necessarily uses $\psi(B_0)$ as its other summand.

Applying (3) now yields

$$|L_\psi(M)| \geq |L_\psi(N)| \geq |U + D| \geq (k - \ell + 1) + (\ell + 1) - 1 = k + 1,$$

as required. The case $\ell = 0$ is included: then D is a singleton. □

Proof of Theorem 1.1. Choose an F-avoiding basis B minimizing $d = |B \setminus A|$. If $d = 0$, there is nothing to prove. Contract $A \cap B$ and delete the elements outside $A \cup B$. The resulting representable matroid has rank d and disjoint bases $A' = A \setminus B$ and $B' = B \setminus A$. Subtract the label of the contracted set from every forbidden label, obtaining a set F' with $|F'| = |F|$.

Every basis of this minor other than B' contains an element of A', so its lift is closer to A than B is. Its label therefore lies in F'. In particular, B' is isolated and the minor has at most $|F| + 1$ distinct basis labels. Theorem 2.2, with the independent set A', gives $d + 1 \leq |F| + 1$. The sharpness example appears in Section 7.3. □

The column operations may change which sets are bases. They preserve the set of basis labels, the independent set I, and the isolation of B, which are precisely the invariants needed for induction. They also preserve representability over the same field; they need not preserve graphicness. No preservation of arbitrary element costs is asserted.

3. The weighted reduction and complementary bases

For a rank-r matroid N on a ground set E_0 of size $2r$, let

$$\mathcal{C}(N) = \{S \subseteq E_0 : S \in \mathcal{B}(N),\ E_0 \setminus S \in \mathcal{B}(N)\}.$$

We call the members of $\mathcal{C}(N)$ *complementary bases*. In the terminology of [10], they are blocks.

Lemma 3.1. *Let $\mathcal{F}$ be a nonempty family of bases of a matroid M. Fix an unconstrained minimum-cost basis A, and choose a minimum-cost member B of $\mathcal{F}$ minimizing* $\operatorname{dist}(A, B)$. *Define*

$$C = A \cap B, \quad N = (M|(A \cup B))/C, \quad A_0 = A \setminus C, \quad B_0 = B \setminus C.$$

Then A_0, B_0 are disjoint bases of N, whose rank is $r = \operatorname{dist}(A, B)$ and whose ground set is $E_0 = A_0 \dot\cup B_0$. Among $S \in \mathcal{C}(N)$, the only one satisfying $C \cup S \in \mathcal{F}$ is B_0.

Proof. Contraction of the independent set C gives the stated minor and bases. Every basis T of N lifts to the basis $C \cup T$ of M. Thus $w(T) \geq w(A_0)$ by optimality of A. For $S \in \mathcal{C}(N)$,

$$w(S) = w(E_0) - w(E_0 \setminus S) \leq w(A_0) + w(B_0) - w(A_0) = w(B_0).$$

If $C \cup S \in \mathcal{F}$, optimality of B forces equality of the two lifted costs. If $S \neq B_0$, the r-element set S contains an element of A_0, whence $\operatorname{dist}(A, C \cup S) = |S \cap B_0| < r$. This contradicts the choice of B. □

The argument of Lemma 3.1 is the strong-block-isolation principle of Liu and Xu [10], applied to general feasibility. In particular, no algebraic assumption is needed for this reduction.

Now suppose N is representable over K. Fix an ordering of E_0 and a full-row-rank representation $U \in K^{r \times 2r}$. Choose $V \in K^{r \times 2r}$ whose rows form a basis of $\ker U$. Then

$$UV^{\mathsf{T}} = 0. \tag{4}$$

The matrix V represents the dual matroid. To see the precise complementary-minor statement, take an r-set S and define $T = E_0 \setminus S$. Projection of $\ker U$ to the coordinates in S has matrix V_S^{T}. Its kernel consists of vectors supported on T that are annihilated by U_T. Both spaces have dimension r, so the projection is invertible exactly when U_T is invertible. Consequently, with all submatrices inheriting the ground-set ordering,

$$c_S = \det(U_S)\det(V_S) \neq 0 \quad \Longleftrightarrow \quad S \in \mathcal{C}(N). \tag{5}$$

This equivalence holds in every characteristic.

4. Polynomial separation and weighted proximity

We first give the elementary identity behind the comparison with exchange families. The basis hypothesis connects the identity to proximity; the cancellation itself is a finite-difference identity.

Lemma 4.1. *Let B be a basis, let $b_1, \dots, b_q \in B$ and $a_1, \dots, a_q \notin B$ be distinct within each list, with $q \geq 1$, and suppose*

$$B(J) = (B \setminus \{b_i : i \in J\}) \cup \{a_i : i \in J\} \quad \text{is a basis for every } J \subseteq [q].$$

Over any field K, every polynomial Q of total degree less than q satisfies

$$\sum_{J \subseteq [q]} (-1)^{|J|} Q(\mathbf{1}_{B(J)}) = 0.$$

Consequently, a polynomial nonzero at exactly one of these basis incidence vectors has degree at least q.

Proof. Substitute $x = \mathbf{1}_B + \sum_{i=1}^{q} t_i(\mathbf{1}_{\{a_i\}} - \mathbf{1}_{\{b_i\}})$ into Q. The resulting polynomial in $t_1, \dots, t_q$ has degree less than q, so each monomial omits some variable t_i. Its evaluations at $t_i = 0$ and $t_i = 1$ cancel in the alternating sum over $\{0,1\}^q$. This remains valid in characteristic two. If exactly one evaluation were nonzero, that sum could not vanish. □

The next identity replaces the cube by the complementary bases of a representable matroid. Its cancellation order is the full rank r, and its coefficients are the complementary-minor products c_S. Its support is exactly the family on which weighted isolation holds.

Lemma 4.2. *With the notation above, every polynomial $Q \in K[x_e : e \in E_0]$ of total degree less than r satisfies*

$$\sum_{\substack{S \subseteq E_0 \\ |S| = r}} c_S Q(\mathbf{1}_S) = 0.$$

Proof. Cauchy–Binet gives

$$H(z) = \det(U \operatorname{diag}(z) V^{\mathsf{T}}) = \sum_{|S| = r} c_S \prod_{e \in S} z_e. \tag{6}$$

This polynomial is homogeneous of degree r, possibly zero. Equation (4) yields the polynomial identity

$$H(\mathbf{1} + z) = \det(UV^{\mathsf{T}} + U \operatorname{diag}(z) V^{\mathsf{T}}) = H(z).$$

Comparing the coefficient of $\prod_{e \in X} z_e$ for $|X| < r$ gives

$$\sum_{\substack{|S| = r \\ S \supseteq X}} c_S = 0. \tag{7}$$

On incidence vectors, $x_e^a = x_e$ for every integer $a \geq 1$. Multilinearizing Q therefore preserves its values and cannot increase its degree. Apply (7) to each resulting monomial. □

Corollary 4.3. *Suppose $\mathcal{C}(N) \neq \varnothing$, and fix $B \in \mathcal{C}(N)$. The minimum degree of a polynomial over K that is nonzero at $\mathbf{1}_B$ and zero at $\mathbf{1}_S$ for every $S \in \mathcal{C}(N) \setminus \{B\}$ is exactly r.*

Proof. A degree less than r contradicts Lemma 4.2, since $c_B \neq 0$. The monomial $\prod_{e \in B} x_e$ gives the upper bound. For $r = 0$, use the constant polynomial 1. □

Proof of Theorem 1.2. Choose B as in Lemma 3.1, with $\mathcal{F} = \{B \in \mathcal{B}(M) : P(\mathbf{1}_B) \neq 0\}$. If $r = \operatorname{dist}(A, B) \leq d$, there is nothing to prove. Otherwise form the minor N in that lemma, which remains K-representable. Substitute $x_e = 1$ for $e \in C$ and $x_e = 0$ for $e \notin A \cup B$. The resulting polynomial Q has degree at most $d < r$ and satisfies $Q(\mathbf{1}_S) = P(\mathbf{1}_{C \cup S})$. By Lemma 3.1, its only nonzero value on complementary bases is at B_0. Lemma 4.2 now gives

$$0 = \sum_{|S| = r} c_S Q(\mathbf{1}_S) = c_{B_0} Q(\mathbf{1}_{B_0}) \neq 0,$$

a contradiction. □

Only the values of P on incidence vectors matter. Accordingly, the degree in Theorem 1.2 may be replaced by the degree of any polynomial with the same values on bases. The theorem does not assert that such a minimum-degree representation can be computed efficiently.

5. The group-algebra determinant

For an abelian group Γ, the group algebra $K[\Gamma]$ has K-basis $\{u_g : g \in \Gamma\}$, with multiplication $u_g u_h = u_{g+h}$. Each element has finite support. The augmentation homomorphism $\varepsilon : K[\Gamma] \to K$ sends every u_g to 1; its kernel I_Γ is the augmentation ideal.

Lemma 5.1. *Let N, U, V be as in Section 3, and let $\psi : E_0 \to \Gamma$. Then*

$$D_\psi = \det(U \operatorname{diag}(u_{\psi(e)}) V^{\mathsf{T}}) = \sum_{|S|=r} c_S u_{\psi(S)} \in I_\Gamma^r. \tag{8}$$

If B is the unique complementary basis with label g, then the coefficient of u_g in D_ψ is $c_B \neq 0$.

Proof. The ring is commutative, so Cauchy–Binet proves the expansion. Applying augmentation to the matrix yields $UV^{\mathsf{T}} = 0$. Every entry therefore belongs to I_Γ, and every product in the determinant has r factors in this ideal. The coefficient claim follows from (5). □

Proposition 5.2. *If $\operatorname{char} K = p$ and $\Gamma = \bigoplus_{j=1}^m \mathbb{Z}_{q_j}$ with $q_j = p^{a_j}$, then*

$$K[\Gamma] \simeq K[t_1, \dots, t_m]/(t_1^{q_1}, \dots, t_m^{q_m}), \qquad I_\Gamma = (t_1, \dots, t_m).$$

In particular, $I_\Gamma^{\nu+1} = 0$, where $\nu = \sum_j (q_j - 1)$.

Proof. Let h_j generate the jth factor and map t_j to $u_{h_j} - 1$. In characteristic p, $(u_{h_j} - 1)^{q_j} = u_{q_j h_j} - 1 = 0$. The resulting homomorphism is surjective, and both algebras have dimension $\prod_j q_j$, so it is an isomorphism. Augmentation sends each t_j to zero. A monomial of total degree at least $\nu + 1$ has an exponent at least q_j for some j, proving nilpotence. □

Proof of Theorem 1.4. Choose an optimal feasible basis B nearest to A and apply Lemma 3.1. Translate T by $-\psi(C)$ after contraction. In the resulting rank-r minor, B_0 is the only complementary basis with a label in the translated allowed set. In particular it is the only one with label $g = \psi(B_0)$. If $r > \nu$, Lemma 5.1 and Proposition 5.2 give $D_\psi = 0$. Its coefficient at u_g is nevertheless $c_{B_0} \neq 0$. This contradiction proves the bound. Sharpness is established in Section 7.3. □

Remark 5.3. Here is an explicit instance of the counterexample in [8, Lemma 6.16, full version]. Identify $E = \{0, \dots, 15\}$ with $\mathbb{F}_2^4$ by binary expansion, label each element by itself, and let $B^* = \{1, \dots, 7\}$. The bases are all seven-element subsets with nonzero label sum, together with B^*. Thus B^* is the unique zero-labeled basis, whereas $\{8, \dots, 14\}$ is a disjoint basis. Their distance is 7, exceeding $D(\mathbb{Z}_2^4) - 1 = 4$.

To see directly why this example falls outside our representability hypothesis, contract $C = \{1, 2, 3\}$ and restrict to the four pairs

$$P_1 = \{4, 5\}, \qquad P_2 = \{6, 7\}, \qquad P_3 = \{8, 9\}, \qquad P_4 = \{12, 13\}.$$

This minor has rank four. Since $\psi(C) = 0$, a four-element set in the minor is a nonbasis precisely when its label sum is zero and it is not $P_1 \cup P_2$. Each pair has label sum 1. A four-element set with distribution $(2, 1, 1)$ among the pairs cannot have zero sum: two elements from different pairs never sum to 1. A set with one element from each pair has sum 6 or 7. Consequently, the only nonbases of size four are the five unions $P_i \cup P_j$ other than $P_1 \cup P_2$. This is the Vámos matroid, which is not representable over any field. The calculation gives an explicit obstruction to representability, independently of Theorem 1.4.

6. Sharp proximity for elementary abelian groups

For $\Gamma = \mathbb{F}_p^m$, nilpotence alone gives the bound $m(p-1)$. To incorporate the number of forbidden labels, we need the minimum support of a nonzero element of I_Γ^r.

Lemma 6.1. *Let K have characteristic p, and let $f \in K[y_1, \dots, y_m]$ be nonzero with degree at most $p-1$ in each variable and total degree at most $m(p-1)-r$, where $0 \le r \le m(p-1)$. Then*

$$|\{y \in \mathbb{F}_p^m : f(y) \neq 0\}| \ge \mu_p(r).$$

Proof. An induction gives the following elementary grid bound. For a nonzero reduced polynomial f, there is a monomial y^β occurring in f such that

$$|\{y \in \mathbb{F}_p^m : f(y) \neq 0\}| \ge \prod_{j=1}^{m} (p - \beta_j). \tag{9}$$

For completeness, write $f = \sum_{i=0}^{t} f_i(y_1, \dots, y_{m-1}) y_m^i$ with $f_t \neq 0$. At each point where f_t is nonzero, the resulting univariate polynomial has at least $p-t$ nonzeros. Apply the induction hypothesis to f_t and append the exponent t. The case of no variables is the nonzero constant polynomial.

For the monomial in (9), define $\alpha_j = p - 1 - \beta_j$. Then $0 \le \alpha_j \le p-1$ and $\sum_j \alpha_j \ge r$. Among all such integer vectors, the product $\prod_j (\alpha_j + 1)$ is minimized with total sum r: reducing positive coordinates can only reduce the product. With the sum fixed, if $1 \le \alpha_i \le \alpha_j < p-1$, replacing (α_i, α_j) by $(\alpha_i - 1, \alpha_j + 1)$ decreases the product, since

$$(\alpha_i + 1)(\alpha_j + 1) - \alpha_i(\alpha_j + 2) = \alpha_j - \alpha_i + 1 > 0.$$

At a minimum, at most one coordinate lies strictly between 0 and $p-1$. Writing $r = a(p-1)+b$, the minimum is therefore $p^a(b+1) = \mu_p(r)$. □

Lemma 6.1 is the usual generalized Reed–Muller minimum-distance bound [4, 9], expressed in complementary degree. The proof also covers coefficient fields larger than $\mathbb{F}_p$.

Lemma 6.2. *Let $\Gamma = \mathbb{F}_p^m$ and $\operatorname{char} K = p$. For $0 \le r \le m(p-1)$, every nonzero $z \in I_\Gamma^r$ satisfies $|\operatorname{supp}(z)| \ge \mu_p(r)$.*

Proof. For $z = \sum_{g \in \Gamma} z_g u_g$, identify its coefficient vector with the function $f_z(g) = z_g$ on $\mathbb{F}_p^m$. The coefficient function of u_0 is represented by

$$f_{u_0}(y) = \prod_{j=1}^{m} (1 - y_j^{p-1}).$$

Multiplication by $t_j = u_{e_j} - 1$ replaces a coefficient function $f(y)$ by $f(y - e_j) - f(y)$. For a polynomial this operation lowers total degree by at least one, unless the result is zero, and does not increase any individual degree. By Proposition 5.2, the ideal I_Γ^r is spanned over K by t^α with $0 \le \alpha_j < p$ and $|\alpha| \ge r$. Its coefficient functions consequently have reduced polynomial representatives of degree at most $m(p-1)-r$. Apply Lemma 6.1. □

Lemma 6.3. *Let $S \subseteq \mathbb{F}_p^m$ be nonempty, let $g \in S$, and let K have characteristic p. There is a polynomial Q over K of degree at most $\lambda_p(|S|-1)$ such that $Q(g) = 1$ and $Q(h) = 0$ for every $h \in S \setminus \{g\}$.*

Proof. Let $d = \lambda_p(|S|-1)$ and $\nu = m(p-1)$. If $d = \nu$, use $Q(y) = \prod_j [1 - (y_j - g_j)^{p-1}]$. Otherwise let $\mathcal{P}_d$ be the space of functions on $\mathbb{F}_p^m$ represented by reduced polynomials of degree at most d. Under the coefficient pairing $\langle z, f \rangle = \sum_h z_h f(h)$, its annihilator in $K[\mathbb{F}_p^m]$ is I_Γ^{d+1}. Indeed, pairing t^α with a polynomial is an iterated finite difference, which vanishes when $|\alpha| > d$. This proves containment; equality follows from dimensions. The reduced monomials of degree at most d are linearly independent as functions on $\mathbb{F}_p^m$, by repeated use of the univariate root bound, and the monomials t^α with $|\alpha| \ge d+1$ form a basis of I_Γ^{d+1}.

If the desired function on S were not a restriction of a member of $\mathcal{P}_d$, finite-dimensional linear algebra would give a vector $(z_h)_{h\in S}$ orthogonal to every such restriction but with $z_g \neq 0$. Extending it by zero gives a nonzero element of I_Γ^{d+1} supported on at most $|S|$ group elements. Lemma 6.2 contradicts this, since $\mu_p(d+1) > |S|$ by the definition of d. □

This separation lemma is also a consequence of classical Reed–Muller duality and minimum distance [4]. It allows us to obtain the sharp elementary-group theorem from polynomial proximity, keeping the same separation argument throughout.

Proof of the upper bound in Theorem 1.3. Fix a minimum-cost F-avoiding basis D and let $g = \psi(D)$. Lemma 6.3 gives a polynomial Q separating g from F, of degree at most $\lambda_p(|F|)$. Substitute the linear forms

$$L_j(x) = \sum_{e\in E} \psi_j(e)x_e$$

to obtain $P(x) = Q(L_1(x), \dots, L_m(x))$ of no larger degree. Every P-feasible basis is F-avoiding, and D is P-feasible. Consequently the minimum cost under $P \neq 0$ equals $w(D)$. Theorem 1.2 supplies such an optimum within $\lambda_p(|F|)$ exchanges of the prescribed unconstrained optimum. The inequality $\lambda_p(k) \leq k$ follows from $\mu_p(r) \geq r+1$. Finally $|F|+1 \leq p^m$ gives $\lambda_p(|F|) \leq m(p-1)$. □

There is an equivalent support interpretation. In the isolated minor, write $F' = F - \psi(C)$ and $g = \psi(B_0)$. The group-algebra determinant is nonzero and supported on $F' \cup \{g\}$. Nilpotence and Lemma 6.2 give

$$\mu_p(r) \leq |\operatorname{supp}(D_\psi)| \leq |F| + 1.$$

The separator and support arguments are dual descriptions of the same obstruction to isolation. The latter also gives the following structural consequence without introducing costs or a forbidden set.

Corollary 6.4. *Let N be a rank-r matroid on $2r$ elements representable in characteristic p, and let $\psi : E(N) \to \mathbb{F}_p^m$. If some complementary basis is the unique complementary basis with its label, then*

$$r \leq m(p-1), \qquad |\{\psi(S) : S \in \mathcal{C}(N)\}| \geq \mu_p(r).$$

Proof. The determinant in Lemma 5.1 is nonzero, and its support is contained in the displayed label set. Apply nilpotence and Lemma 6.2. □

7. Consequences, algorithms, and sharpness

7.1. **Additive labels and simultaneous constraints.** The polynomial theorem gives a short proof of the weaker $|F|$ bound and covers labels over characteristic-zero fields as well.

Corollary 7.1. *Let M be K-representable, let $\psi : E \to K^m$, and let $F \subseteq K^m$ be finite. If an F-avoiding basis exists, every unconstrained minimum-cost basis lies within $|F|$ exchanges of some minimum-cost F-avoiding basis.*

Proof. Choose a minimum-cost F-avoiding basis B and let $g = \psi(B)$. For each $f \in F$, choose a coordinate $j(f)$ with $g_{j(f)} \neq f_{j(f)}$. Define

$$L_j(x) = \sum_{e\in E} \psi_j(e)x_e, \qquad P(x) = \prod_{f\in F} (L_{j(f)}(x) - f_{j(f)}).$$

Then $\deg P \leq |F|$, $P(\mathbf{1}_B) \neq 0$, and P vanishes at every basis with label in F. All P-feasible bases are therefore F-avoiding, and their optimum has the same cost as B. Theorem 1.2 gives the conclusion. □

Corollary 7.2. *Corollary 7.1 holds for labels in any torsion-free abelian group if M is representable over a field of characteristic zero.*

Proof. The subgroup generated by the finitely many element labels and forbidden labels is a finitely generated torsion-free abelian group, hence isomorphic to $\mathbb{Z}^s$. Embed its coordinates in the prime subfield of K and apply Corollary 7.1. □

For comparison, the unweighted $|F|$ bound for ordered abelian groups is already known for arbitrary matroids [8]. The conclusion above includes real-cost optimality but assumes characteristic-zero representability.

Corollary 7.3. *Suppose M is representable in characteristic p. For $t = 1, \dots, \ell$, let $\psi_t : E \to \mathbb{F}_p^{m_t}$ and $F_t \subseteq \mathbb{F}_p^{m_t}$. If a basis avoiding all the F_t exists, then every unconstrained optimum lies within*

$$\sum_{t=1}^{\ell} \lambda_p(|F_t|)$$

exchanges of a minimum-cost simultaneously feasible basis. In particular, ℓ nonzero-label constraints require at most ℓ exchanges.

Proof. Fix an optimal simultaneously feasible basis B and let $g_t = \psi_t(B)$. Apply Lemma 6.3 to $F_t \cup \{g_t\}$. Substitute $L_{tj}(x) = \sum_e \psi_{tj}(e)x_e$ into its separating polynomial to obtain P_t of degree at most $\lambda_p(|F_t|)$. The product $\prod_t P_t$ is nonzero at B, and its nonvanishing implies simultaneous feasibility. Its optimum has the cost of B, so Theorem 1.2 applies with the sum of the degrees. □

The common characteristic is essential to this argument. The proof uses one product of separating polynomials, rather than successive applications of single-constraint proximity.

7.2. **A direct exact algorithm for additive costs.** Let $n = |E|$ and $\rho = \operatorname{rank}(M)$. First find an unconstrained optimum A by the greedy algorithm. If a proximity bound d applies, enumerate all sets $(A \setminus X) \cup Y$ with $X \subseteq A$, $Y \subseteq E \setminus A$, and $|X| = |Y| \le d$. Test whether each set is a feasible basis and retain the cheapest. The number of candidates is

$$\sum_{i=0}^{\min\{d,\rho,n-\rho\}} \binom{\rho}{i}\binom{n-\rho}{i} \le (d+1)\max\{1,n\}^{2d}. \tag{10}$$

If no candidate is feasible, the proximity theorem certifies infeasibility. This is an oracle-operation bound; the bit complexity also includes the costs of arithmetic, label operations, and evaluating the constraint. For rational costs, all comparisons can be exact.

For polynomial feasibility use $d = \deg P$. For elementary-group avoidance use $d = \lambda_p(|F|)$, giving $n^{O((p-1)\log_p(|F|+1))}$ candidate checks for $|F| \ge 1$. In the binary case this is $n^{O(\log(|F|+1))}$. These are direct enumeration bounds, not fixed-parameter algorithms. The separating polynomial in Corollary 7.1 need not be known: group feasibility is tested directly.

For arbitrary abelian labels on a representable matroid, Theorem 1.1 gives an $n^{O(|F|)}$ feasibility algorithm: start from any basis and search its radius-$|F|$ neighborhood. This conclusion does not assert that a cheapest feasible basis lies in that neighborhood.

7.3. **Sharp examples.** Each construction uses a path in which every edge is replaced by a parallel pair. A spanning tree chooses one edge from each pair. Call the choices a_i, b_i, give a_i cost zero and b_i cost one, and let A choose all the a_i. Thus A is the unique unconstrained optimum.

The unweighted bound. With k pairs, label every a_i by 0 and every b_i by 1 in $\mathbb{Z}$, and take $F = \{0, \dots, k-1\}$. The unique avoiding tree chooses all b_i and is at distance $k = |F|$ from A. For $k = 0$, use the one-vertex graph. This proves sharpness in Theorem 1.1 over every representation field.

Polynomial degree. With d pairs, let $P(x) = \prod_{i=1}^{d} x_{b_i}$. The unique feasible basis chooses every b_i and has distance d from A. This proves sharpness in Theorem 1.2.

The exact elementary-group bound. Fix $0 \le k < p^m$, let $r = \lambda_p(k)$, and write $r = a(p-1)+b$ with $0 \le b < p-1$. For each $j \le a$, take $p-1$ pairs with labels $\psi(a_i) = e_j$ and $\psi(b_i) = 0$. If $b > 0$, take b further pairs with nonzero label e_{a+1}. The realized tree labels form the box

$$\mathbb{F}_p^a \times \{0, \dots, b\} \times \{0\}^{m-a-1}$$

when $b > 0$, and $\mathbb{F}_p^a \times \{0\}^{m-a}$ when $b = 0$. In either case there are $\mu_p(r)$ labels. Only the all-b_i tree has label zero, since each coordinate count lies between 0 and $p-1$.

Forbid all nonzero realized labels, and add arbitrary unrealized labels until $|F| = k$. This is possible because $\mu_p(r) - 1 \le k \le p^m - 1$. The only feasible tree is the all-b_i tree, at distance r from A. The case $r = 0$ is the one-vertex graph. This proves the sharpness assertion of Theorem 1.3.

The p-group bound. For each cyclic generator h_j of order $q_j = p^{a_j}$, take $q_j - 1$ pairs with labels $h_j, 0$. Let the sole admissible label be zero. No coordinate count can wrap around its modulus, so the only feasible tree chooses all zero-labeled edges. Its distance from A is $\sum_j (q_j - 1) = \nu(\Gamma)$, proving sharpness in Theorem 1.4.

8. Realizable valuated costs

The vanishing-moment identity also applies to certain nonadditive basis costs. The role of cost comparison is then played by the non-Archimedean rule: a finite sum with a unique term of smallest valuation cannot vanish.

Let L be a field with a real-valued non-Archimedean valuation $v : L \to \mathbb{R} \cup \{\infty\}$, with $v(0) = \infty$. Let $X \in L^{\rho \times E}$ have full row rank and represent M. For arbitrary $w : E \to \mathbb{R}$, define

$$q(B) = v(\det X_B) + w(B), \qquad B \in \mathcal{B}(M). \tag{11}$$

We refer to these as realizable valuated costs with a linear perturbation. The hypothesis concerns the costs themselves, not just representability of the underlying matroid.

Theorem 8.1. *Let q be as in (11), and let $P \in L[x_e : e \in E]$ have degree at most d and be nonzero at some basis incidence vector. For every minimizer A of q, there is a minimizer B of*

$$h(B) = q(B) + v(P(\mathbf{1}_B))$$

with $\operatorname{dist}(A, B) \le d$. Here $h(B) = \infty$ when $P(\mathbf{1}_B) = 0$.

Proof. Choose an h-minimizer B nearest to A and suppose $r = \operatorname{dist}(A, B) > d$. Let $C = A \cap B$, and form the minor on $E_0 = (A \cup B) \setminus C$ with disjoint bases A_0, B_0 as before. Normalize the independent columns C by invertible row operations and delete their rows and columns. This gives a representation U of the minor for which

$$\det X_{C \cup S} = \epsilon_S \alpha \det U_S$$

for every r-set S, with a fixed $\alpha \neq 0$ and signs $\epsilon_S \in \{1, -1\}$. Consequently $q'(S) = v(\det U_S) + w(S)$ differs from $q(C \cup S)$ by a common constant. After the usual substitutions in P, write the resulting polynomial as Q and define $h'(S) = q'(S) + v(Q(\mathbf{1}_S))$. Thus A_0 minimizes q' and B_0 minimizes h'.

Choose a kernel-row matrix V for U and define $c_S = \det U_S \det V_S$. Complementary minors satisfy

$$\det V_S = \eta_S \beta \det U_{E_0 \setminus S}$$

for a fixed $\beta \neq 0$ and signs η_S. For example, after normalizing $U = [I \; W]$, this follows by taking $V = [-W^{\mathsf{T}} \; I]$ and expanding the minors. For each nonzero term below, it follows that

$$v(c_S Q(\mathbf{1}_S)) = h'(S) + q'(E_0 \setminus S) + v(\beta) - w(E_0). \tag{12}$$

Lemma 4.2, applied over L, gives $\sum_{|S|=r} c_S Q(\mathbf{1}_S) = 0$. Its B_0 term is nonzero. Every other nonzero term has $q'(E_0 \setminus S) \ge q'(A_0)$ and $h'(S) > h'(B_0)$: equality in the latter would give an h-minimizer $C \cup S$ closer to A. By (12), the B_0 term has uniquely smallest valuation. This contradicts the non-Archimedean rule and proves $r \le d$. □

The penalty $v(P(\mathbf{1}_B))$ disappears when all nonzero evaluations have the same valuation. In particular this holds for polynomials over a subfield on which v is trivial. The following observation removes the coefficient restriction in equal characteristic.

Corollary 8.2. *Suppose v is trivial on the prime subfield k of L. For every polynomial $P \in L[x_e : e \in E]$ of degree at most d with a feasible basis, every q-minimizer lies within d exchanges of a q-minimum basis satisfying $P(\mathbf{1}_B) \neq 0$.*

Proof. Choose a k-basis $a_1, \dots, a_s$ of the span of the coefficients of P, and write $P = \sum_i a_i P_i$, with $P_i \in k[x]$ of degree at most d. At incidence vectors, P vanishes exactly when all P_i vanish. Choose a q-optimal P-feasible basis D and an i with $P_i(\mathbf{1}_D) \neq 0$. The P_i-feasible bases form a subset of the P-feasible bases containing D, so their optimum is $q(D)$. All nonzero evaluations of P_i have valuation zero. Apply Theorem 8.1 to P_i. □

The hypothesis holds in positive characteristic and in residue characteristic zero. We make no such pure-feasibility assertion for arbitrary coefficients in mixed characteristic; the penalty theorem still applies there.

Corollary 8.3. *If $\operatorname{char} L = p$ and $\psi : E \to \mathbb{F}_p^m$, every minimizer of q lies within $\lambda_p(|F|)$ exchanges of a q-minimum F-avoiding basis, whenever one exists. This bound is sharp.*

Proof. Choose a q-optimal F-avoiding basis D and let $g = \psi(D)$. Use Lemma 6.3 over $\mathbb{F}_p$ to separate g from F, then substitute the linear label forms. The resulting polynomial has degree at most $\lambda_p(|F|)$; its feasible bases are F-avoiding and include D. Apply Corollary 8.2. Constant representations with linear costs include the sharp examples in Section 7.3. □

These objectives genuinely extend additive costs. Over $L = K((t))$ with its t-adic valuation, take $w = 0$ and

$$X = \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1+t \end{pmatrix}.$$

All six pairs are bases; $q(\{3,4\}) = 1$ and the other five costs are zero. An additive objective, even up to a common constant, would give equal values to the three sums of costs of complementary pairs. Here those sums are $1, 0, 0$. Our conclusions concern the realizable costs in (11); they do not assert proximity for arbitrary abstract valuated matroids or a greedy algorithm for minimizing q.

9. Toward weighted proximity for arbitrary abelian labels

The unweighted bound of Chen [3], reproved in Theorem 1.1, applies to every representable matroid. The weighted question for arbitrary abelian labels requires more: the reduction isolates a feasible optimum only among complementary bases. The column operations of Section 2 do not preserve this family or its costs. We formulate a sufficient condition for a weighted linear bound, already of interest for graphic matroids, and explain why augmentation-ideal membership alone cannot prove it. The discussion concerns what our methods establish, rather than a claim that the weighted proximity conjecture remains open: Liu, Oki, and Xu have communicated a resolution, as recorded in the closing note. For group-independent bounds, it suffices to consider finite abelian groups.

Proposition 9.1. *Any group-independent weighted proximity bound valid for every finite abelian labeling of a fixed class of finite matroids is valid for arbitrary abelian labelings of that class.*

Proof. Let H be the subgroup generated by the element labels and the finite forbidden set. Then $H \simeq \mathbb{Z}^s \oplus T$ for a finite abelian group T. The set consisting of all basis labels and all forbidden labels is finite. Choose a positive integer q larger than the absolute values of all nonzero coordinate differences between members of this set in the free part. Reduction modulo q on the free part, with the identity on T, is injective on that finite set. It therefore preserves the feasible bases exactly. The costs and distances are unchanged. □

Question 9.2. Let N be a rank-r graphic matroid on $2r$ elements, and let $\psi : E(N) \to \Gamma$ label its elements in an arbitrary abelian group. Suppose $B \in \mathcal{C}(N)$ is the only complementary basis with label $\psi(B)$. Must

$$|\{\psi(S) : S \in \mathcal{C}(N)\}| \geq r + 1$$

hold? More weakly, is this cardinality at least $1 + r/C$ for an absolute constant C?

Proposition 9.3. *An affirmative answer to the weak form of Question 9.2 implies the weighted $C|F|$ proximity bound for graphic matroids. The sharp form implies the weighted $|F|$ bound.*

Proof. Apply Lemma 3.1 to a nearest optimal F-avoiding basis. Graphic matroids are closed under minors. In the reduced minor, B_0 has a unique label among complementary bases and all other complementary labels lie in the translated forbidden set. Thus the label set has size at most $|F| + 1$. The proposed lower bound gives $1 + r/C \leq |F| + 1$, so $r \leq C|F|$. □

For elementary abelian labels, Corollary 6.4 proves the stronger lower bound $\mu_p(r)$. For torsion-free labels on a characteristic-zero representable matroid, the linear-factor separator in Corollary 7.1, together with Corollary 4.3, proves the lower bound $r + 1$. The cases of this structural question not covered by these arguments involve torsion beyond a single elementary abelian group.

9.1. **Why ideal membership alone is insufficient.** One might hope to prove Question 9.2 by extracting a support bound from $D_\psi \in I_\Gamma^r$. The following two examples show why more information about the determinant is required.

Proposition 9.4. *Let $a \geq 2$ and $R = \mathbb{F}_p[\mathbb{Z}_{p^a}]$. For a generator u, the nonzero element*

$$1 - u^{p^{a-1}} = (1 - u)^{p^{a-1}}$$

has support two and belongs to $I_\Gamma^{p^{a-1}}$.

Proof. The identity is the Frobenius identity in characteristic p. The group elements 0 and p^{a-1} are distinct modulo p^a. □

Even working integrally does not remove the mixed-prime obstruction.

Proposition 9.5. *In $R = \mathbb{Z}[\mathbb{Z}_6] \simeq \mathbb{Z}[x]/(x^6 - 1)$, there is a nonzero element supported on four group elements, with all nonzero coefficients equal to 1 or -1, that belongs to every positive power of the augmentation ideal.*

Proof. Define

$$q = x^4 + x^3 - x - 1 = (x - 1)(x + 1)(x^2 + x + 1).$$

This element is nonzero in R and has the asserted coefficients. Since $q(x^2 - x + 1) = x^6 - 1 = 0$ in R and $x^2 - x + 1 = 1 + x(x - 1)$, we have $q = -x(x - 1)q$. Iterating gives

$$q = (-x)^s(x - 1)^s q \in I_\Gamma^s \qquad (s \geq 1),$$

as $I_\Gamma = (x - 1)$. □

Neither proposition is a counterexample to matroid proximity. In particular, we do not claim that these elements arise from a matroid determinant with an isolated complementary label. They show that ideal membership and a nonzero unit coefficient do not by themselves control support.

9.2. **The additional determinant structure.** In a reduced minor with disjoint bases A_0, B_0, a representation can be normalized to $U = [I\ W]$ with W invertible. We may choose $V = [-W^{\mathsf{T}}\ I]$. The determinant in (8) then has the explicit form

$$\det\left(W\,\mathrm{diag}(u_{\psi(b_j)}) - \mathrm{diag}(u_{\psi(a_i)})W\right). \tag{13}$$

For a regular matroid, W can be chosen totally unimodular. The remaining step within this determinant approach is therefore to use the special minors and cancellations in (13), together with isolation of one complementary label. A support theorem for arbitrary elements of the augmentation ideal cannot supply this missing step.

For elementary groups, isolation and the support bound complete this strategy. The step not supplied by our argument is to exploit the additional determinant structure for non-elementary torsion.

For nonrepresentable matroids, even the unweighted conjecture remains outside the present argument. An extension of Theorem 2.2 to all matroids would give another route to the unweighted bound by the same nearest-basis reduction. The missing ingredient is an alternative to the label-preserving column operations used in its proof.

Appendix A. Compatible exchange families for spanning trees

Theorem 1.1 gives sharp proximity without constructing large compatible exchange families. We record a separate structural property of spanning trees, which gives such a family explicitly.

Throughout this section, B and A are edge-disjoint spanning trees on the same vertex set, with r edges each. Edges retain their identities, including parallel edges.

Definition A.1. A set $S \subseteq A$ is *compatible with* B if there is an injection $b : S \to B$ such that

$$(B \setminus b(J)) \cup J \quad \text{is a spanning tree for every } J \subseteq S.$$

No condition is imposed on the complementary edge set. The injections for different compatible families may use the same edges of B.

We first record a simple way to verify compatibility. Let $P_B(a)$ be the unique path in B between the endpoints of $a \notin B$.

Lemma A.2. *Let $a_1, \dots, a_q$ be distinct edges outside B, and let $b_1, \dots, b_q$ be distinct edges of B. Suppose*

$$b_i \in P_B(a_i), \qquad b_i \notin P_B(a_j) \quad \text{whenever } i < j.$$

Then these pairs form a compatible exchange family.

Proof. Fix any subset of the pairs and perform its exchanges in increasing index order. At the exchange of a_j for b_j, the original path $P_B(a_j)$ is still present: it contains none of the previously removed edges. Since the current graph is a tree, adding a_j creates precisely the cycle $P_B(a_j) \cup \{a_j\}$. Removing b_j therefore leaves a tree. This proves the assertion for every subset. □

Theorem A.3. *The edge set A can be partitioned into three sets, each compatible with B. In particular, there is a compatible exchange family of size at least $\lceil r/3 \rceil$.*

Proof. Root B at a vertex o. Independently root A at any vertex and orient its edges away from that root. Thus every vertex is the head of at most one oriented edge of A.

Choose an ordering of the children at each vertex of B. Let $\prec_1$ be the resulting depth-first *preorder*: a vertex is visited before its children, and each child subtree is completed before the next is visited. Let $\prec_2$ be the preorder obtained by reversing the order of the children at every vertex. Both orders place ancestors before descendants. For two incomparable vertices, their order is reversed: at their lowest common ancestor, the two child subtrees containing them are visited in opposite orders.

For an oriented edge $a = (u \to v)$ of A, distinguish the case where v is a strict ancestor of u in B. Assign all such edges to U. Of the remaining edges, assign those with $u \prec_1 v$ to F_1

and the rest to F_2. Every edge in F_2 has $u \prec_2 v$: its endpoints must be incomparable, since an edge from an ancestor to a descendant is forward in both orders. This partitions A into U, F_1, F_2.

The forward families. For an edge $(u \to v) \in F_i$, pair it with the parent edge of v in B. The vertex v is not o, because o comes first in $\prec_i$. These parent edges are distinct, since the heads of the oriented A-edges are distinct. Order the pairs by decreasing position of their heads in $\prec_i$. For $(u \to v) \in F_i$, the vertex u is outside the subtree rooted at v: all proper descendants of v occur after v in preorder. Thus $P_B(uv)$ contains the parent edge of v. Every vertex on this path is an ancestor of u or of v, and therefore occurs no later than v in $\prec_i$. The path cannot contain the parent edge paired with an earlier pair, whose head occurs strictly later than v. Lemma A.2 proves compatibility of both forward families.

The ancestor family. For $(u \to v) \in U$, pair it with the first edge of the B-path from v down to u. This edge has upper endpoint v. The paired B-edges are distinct because the heads v are distinct. Order these pairs by nondecreasing depth of their heads in B. The path of an edge with head v_j lies entirely in the subtree rooted at v_j. It cannot contain the paired edge with upper endpoint v_i for any earlier index i: either v_i is shallower, or it is a different vertex at the same depth. Each path contains its own paired edge. Lemma A.2 therefore proves compatibility of U.

One of these three compatible sets has at least $\lceil r/3 \rceil$ edges, as required. □

Contracting common edges extends the size conclusion to arbitrary tree pairs, with r equal to their exchange distance. To compare with [8, Section 5.1.3, full version], their elementary k-exchange property requires k-element sets $X \subseteq B_1 \setminus B_2$ and $Y \subseteq B_2 \setminus B_1$ and a bijection $\varphi : X \to Y$ such that $(B_1 \setminus Z) \cup \varphi(Z)$ is a basis for every $Z \subseteq X$. This definition is one-sided: it imposes no complementary exchange condition on B_2.

For a k-element compatible set $S \subseteq A \setminus B$ with injection $b : S \to B \setminus A$, take $B_1 = B$, $B_2 = A$, $X = b(S)$, $Y = S$, and $\varphi = b^{-1}$. For each $Z \subseteq X$, compatibility applied to $J = b^{-1}(Z)$ gives precisely the required basis. Thus graphic matroids are elementarily $(3k-2, k)$-weakly base orderable for every $k \geq 1$: if $r \geq 3k-2$, then $\lceil r/3 \rceil \geq k$. This improves the $(3k^3, k)$ parameter of [8, Theorem 5.27, full version].

To relate this structure to isolation, suppose B is the only tree with its label and (b_i, a_i), $1 \leq i \leq q$, are compatible pairs. Define $\delta_i = \psi(a_i) - \psi(b_i)$. No nonempty subset of these increments sums to zero. Thus the $q+1$ prefix sums are pairwise distinct and give $q+1$ different tree labels. Together with the nearest-basis reduction, the theorem alone gives the weaker bound $3|F|$ for graphic matroids. Theorem 1.1 improves this to $|F|$ by working in the larger class of representable matroids.

Compatibility is one-sided: the complementary edges need not remain a tree. For example, take $B = \{01, 12, 23\}$ and $A = \{02, 03, 13\}$ in K_4. The pairs $(12, 02)$ and $(23, 03)$ are compatible from B. Performing both gives the star $\{01, 02, 03\}$, whereas the complementary edges $\{12, 13, 23\}$ form a triangle and isolate vertex 0. Thus these families do not directly address the complementary-basis isolation used for costs.

Note added

In personal communication in October 2026, Siyue Liu informed the author that they along with co-authors Taihei Oki, and Chao Xu have obtained a resolution of the weighted proximity conjecture for group-labeled matroid bases. Their manuscript was not available to the author when this note was prepared, so no comparison of their methods or precise results with those proved here is attempted. References above to limitations of our arguments should be read in light of this communication. The author thanks Siyue Liu for sharing this information.

AI-Use Disclosure

The appendix gives an elementary combinatorial proof of a linear (3|F|) proximity bound for spanning trees with arbitrary abelian labels which was obtained independently. During

the proof verification and literature exploration using AI, ChatGPT provided the insight for the alternative proof of the (|F|) bound presented here; the bound itself was independently established by Chen [3]. The author takes full responsibility for the correctness, originality, and integrity of the final manuscript.

References

1. Noga Alon and Zoltán Füredi, *Covering the cube by affine hyperplanes*, European Journal of Combinatorics **14** (1993), no. 2, 79–83.
2. Kristóf Bérczi and Tamás Schwarcz, *Partitioning into common independent sets via relaxing strongly base orderability*, Journal of Combinatorial Theory, Series A **202** (2024), 105817.
3. Lizhong Chen, *Sharp proximity bounds for group-labelled matroid bases*, Zenodo preprint, 2026, 22 September 2026. https://doi.org/10.5281/zenodo.22885939.
4. Philippe Delsarte, Jean-Marie Goethals, and F. Jessie MacWilliams, *On generalized Reed–Muller codes and their relatives*, Information and Control **16** (1970), no. 5, 403–442.
5. Friedrich Eisenbrand, Lars Rohwedder, and Karol Węgrzycki, *Sensitivity, proximity and FPT algorithms for exact matroid problems*, 65th IEEE Annual Symposium on Foundations of Computer Science, IEEE, 2024, pp. 1610–1620.
6. Dániel Garamvölgyi, Ryuhei Mizutani, Taihei Oki, Tamás Schwarcz, and Yutaro Yamaguchi, *Towards the proximity conjecture on group-labeled matroids*, 52nd International Colloquium on Automata, Languages, and Programming, Leibniz International Proceedings in Informatics, vol. 334, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025, pp. 85:1–85:17.
7. Yahya Ould Hamidoune, *Extensions of the Scherk–Kemperman theorem*, Journal of Combinatorial Theory, Series A **117** (2010), no. 7, 974–980.
8. Florian Hörsch, András Imolay, Ryuhei Mizutani, Taihei Oki, and Tamás Schwarcz, *Problems on group-labeled matroid bases*, 51st International Colloquium on Automata, Languages, and Programming, Leibniz International Proceedings in Informatics, vol. 297, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024, pp. 86:1–86:20.
9. Elodie Leducq, *A new proof of Delsarte, Goethals and Mac Williams theorem on minimal weight codewords of generalized Reed–Muller codes*, Finite Fields and Their Applications **18** (2012), no. 3, 581–586.
10. Siyue Liu and Chao Xu, *On the congruency-constrained matroid base*, Integer Programming and Combinatorial Optimization, Lecture Notes in Computer Science, vol. 14679, Springer, 2024, pp. 280–293.
11. John E. Olson, *A combinatorial problem on finite abelian groups, I*, Journal of Number Theory **1** (1969), no. 1, 8–10.

Email address: abhardwaj@iitb.ac.in

Department of Industrial Engineering and Operations Research,, Indian Institute of Technology Bombay,, Mumbai, India 400076