Lorentz Control in Additive Combinatorics: Tail-Free L^3 Estimates and Weighted Propagation

Liz Lemma   Future Detective

January 16, 2026

Abstract

Bloom introduced an L^3 'control' parameter $\kappa(A)$ defined by the uniform bound $\sum_x (1_A * 1_B(x))^3 \leq \kappa(A)|A|^2|B|^2$ for all finite B, and showed that improvements in control-to-structure inequalities propagate to convex-set growth, sum-product, and Balog–Szemerédi–Gowers-type results. A technical step in the paper replaces indicators 1_B by general functions using a dyadic layer-cake argument, at the cost of a large tail term (of size κ^{100}) that is harmless for the paper's exponent-level results but obstructs weighted/asymmetric variants.

We identify the correct functional-analytic formulation: the convolution operator $T_A : f \mapsto 1_A * f$ has a sharp strong-type bound from the discrete Lorentz space $\ell^{3/2,1}(G)$ to $\ell^3(G)$. We define the Lorentz control constant $\mathsf{C}(A)$ as the operator norm of T_A on $\ell^{3/2,1}$ and prove the exact identity $\mathsf{C}(A) = \kappa(A)^{1/3}$. This yields a tail-free extension principle: for all finitely supported $f \geq 0$, $\|1_A * f\|_3 \leq \kappa(A)^{1/3}|A|^{2/3}\|f\|_{3/2,1}$.

We then repackage Bloom's symmetry-set arguments in a weighted form, replacing repeated dyadic truncations by Lorentz-norm bookkeeping. As a result, all of Bloom's 'propagation' applications extend cleanly to weighted/asymmetric settings (including continuous weights and measures), and one can isolate exactly where polylogarithmic losses arise when working in $\ell^{3/2}$ instead of $\ell^{3/2,1}$.

Table of Contents

1. 1. Introduction and motivation: control as an operator norm; why Lorentz spaces remove tail losses; overview of applications (weighted/asymmetric, stability, portability).

2. 2. Preliminaries: convolutions on abelian groups; rearrangements; discrete Lorentz spaces $\ell^{p,1}$; basic inequalities (triangle, Hölder in Lorentz form) used later.

3. 3. Lorentz control constant $\mathsf{C}(A)$: definition; basic properties (scaling, monotonicity, disjoint union subadditivity); evaluation on indicators.

4. 4. Equivalence theorem $\mathsf{C}(A) = \kappa(A)^{1/3}$: (a) indicator control implies Lorentz bound; (b) Lorentz bound implies indicator control; discussion of when polylog losses appear (only when converting $\ell^{3/2,1} \leftrightarrow \ell^{3/2}$).

5. 5. Tail-free function extension principle: replace Bloom's Lemma 2 with a Lorentz-strong estimate; derive clean corollaries for dyadic layer sets without κ^{100} terms.

6. 6. Weighted symmetry-set machinery: define weighted symmetry sets via thresholds on $1_A \circ w$; prove Lorentz versions of the key bounds on $\|1_A \circ 1_S\|_{3/2}$ and the auxiliary inner-product inequalities.

7. 7. Propagation results in weighted/asymmetric form: restate Bloom's main applications (energy, sum/difference sets, BSG extraction, convex-function decomposition) with weights/measures and Lorentz control.

8. 8. Examples and sharpness: convex sets (Szemerédi–Trotter) yield $\mathsf{C}(A) \lesssim |A|^{-1/3}$; unions (subadditivity); toy weighted examples where Lorentz beats $\ell^{3/2}$; discussion of limitations (cannot by itself improve sumset exponent).

9. 9. Outlook (2026 directions): finite-field L^4 analogues; automated inequality search with Lorentz primitives; potential stability/inverse theory for Lorentz control.

1 1. Introduction and motivation: control as an operator norm; why Lorentz spaces remove tail losses; overview of applications (weighted/asymmetric, stability, portability).

We recast Bloom's L^3-*control* hypothesis as an operator norm statement for the convolution operator $T_A\colon f \mapsto 1_A * f$. The point is that the quantitative input in Bloom's argument is not tied to indicators 1_B as such; rather, it is a bound on the third moment of representation functions, and such bounds are more naturally expressed as a mapping property between function spaces. Once the correct domain space is chosen, one obtains a formulation that is simultaneously (i) equivalent to Bloom's original invariant when tested on indicators, and (ii) stable under the layer-cake manipulations that appear throughout the propagation and extraction steps.

The guiding observation is that the inequality

$$\sum_x (1_A * 1_B(x))^3 \;\lesssim\; |A|^2\,|B|^2$$

is a priori an estimate for $\|1_A * 1_B\|_{\ell^3} = \left(\sum_x (1_A * 1_B(x))^3\right)^{1/3}$, and so it should be compared to $|A|^{2/3}|B|^{2/3}$. In other words, the natural normalization suggests that T_A ought to map some version of $\ell^{3/2}$ to ℓ^3 with operator norm $\asymp |A|^{2/3}$ times a parameter depending only on A. The exponent $3/2$ is forced by scaling (or, in the discrete setting, by homogeneity): Hölder's inequality indicates that ℓ^3 is dual to $\ell^{3/2}$, and the third moment is the quantity that propagates in Bloom's symmetry-set arguments.

However, Bloom's paper requires more than the indicator-to-indicator inequality. At several points one needs to replace 1_B by a general nonnegative function f, interpreted as a weighted multiset, and to control $\|1_A * f\|_3$ in terms of a simple functional of f. The naive choice $\|f\|_{3/2}$ is not sufficiently well adapted to dyadic decompositions: if one writes f as a sum of dyadic pieces,

$$f \approx \sum_{k\in\mathbb{Z}} 2^k 1_{E_k}, \qquad E_k := \{x : 2^k \le f(x) < 2^{k+1}\},$$

then the triangle inequality in ℓ^3 yields

$$\|1_A * f\|_3 \;\le\; \sum_k 2^k\,\|1_A * 1_{E_k}\|_3.$$

If one inserts Bloom's control bound for each E_k, one obtains a sum of the form $\sum_k 2^k |E_k|^{2/3}$. This sum is not $\|f\|_{3/2}$; it is the endpoint Lorentz quantity $\|f\|_{3/2,1}$. Any attempt to estimate $\sum_k 2^k |E_k|^{2/3}$ by $\|f\|_{3/2}$ introduces either a logarithmic factor (when f has bounded range) or a truncation/tail term (when f has many scales). In Bloom's implementation, such tail terms are suppressed by additional hypotheses and then absorbed into a large power of κ, producing the familiar κ^{100}-type losses.

The Lorentz space $\ell^{3/2,1}(G)$ is exactly the space designed to measure the layer-cake sum without loss. Concretely, for finitely supported $f \geq 0$, the quasi-norm $\|f\|_{3/2,1}$ is equivalent (up to absolute constants) to the dyadic expression

$$\sum_{k \in \mathbb{Z}} 2^k \, |E_k|^{2/3}.$$

Thus the triangle inequality in ℓ^3, combined with control estimates on indicators, produces a clean and tail-free extension to general weights:

$$\|1_A * f\|_3 \;\lesssim\; (\text{control parameter for } A)^{1/3} \, |A|^{2/3} \, \|f\|_{3/2,1}.$$

In this formulation there is no need to discard small values of f or to impose an *ad hoc* cutoff to guarantee a bounded number of dyadic layers. Any later conversion from $\|f\|_{3/2,1}$ to $\|f\|_{3/2}$ is then clearly separated as a secondary step, and any polylogarithmic loss is attributable solely to this conversion rather than to the control mechanism itself.

The second key point is that this Lorentz-operator formulation is not merely comparable to Bloom's original invariant; it is essentially equivalent. Indeed, when $f = 1_B$ is an indicator, the Lorentz quasi-norm satisfies $\|1_B\|_{3/2,1} \asymp |B|^{2/3}$. Consequently, an inequality of the form

$$\|1_A * f\|_3 \;\leq\; \mathsf{C}(A) \, |A|^{2/3} \, \|f\|_{3/2,1}$$

immediately implies

$$\|1_A * 1_B\|_3^3 \;\leq\; \mathsf{C}(A)^3 \, |A|^2 \, |B|^2$$

for all finite B, which is Bloom's control condition with $\kappa(A) = \mathsf{C}(A)^3$ (up to the normalization constant in $\|1_B\|_{3/2,1}$). Conversely, assuming the indicator control for all B, the dyadic decomposition argument described above gives the weighted inequality and hence bounds $\mathsf{C}(A)$ by $\kappa(A)^{1/3}$. This exactness matters in applications: it means that any improvement in the control parameter (coming, for instance, from incidence geometry, sum-product input, or structural hypotheses on A) transfers without degradation to the operator norm $\mathsf{C}(A)$, and conversely any operator-norm bound yields the corresponding moment inequality for indicators.

From the perspective of applications, the operator-norm package is useful for three reasons.

First, it is inherently *weighted and asymmetric.* Many arguments in additive combinatorics now pass through weighted models: sampling measures, multiplicities, entropy-type weights, or intermediate functions produced by Cauchy–Schwarz. Once we know that T_A maps $\ell^{3/2,1}$ to ℓ^3 with norm $\mathsf{C}(A)|A|^{2/3}$, any such weight f may be substituted directly, with no need to return to an indicator by a dyadic pigeonhole or to discard a portion of the mass. Likewise, asymmetric variants (e.g. $1_A * f$ with f supported on a different set or having different scale properties) fit the same formalism.

Second, the Lorentz formulation is *stable under decomposition.* If A is partitioned into disjoint pieces $A = \bigsqcup_i A_i$, then $1_A * f = \sum_i 1_{A_i} * f$, and the triangle inequality in ℓ^3 yields subadditivity of the corresponding operator norms. This is precisely the kind of bookkeeping needed in iterative pruning and energy increment arguments, where one repeatedly removes structured or unstructured parts and must track how the control parameter evolves. Expressing control as $\mathsf{C}(A)$ makes such steps transparent, while keeping the dependence on constants free of auxiliary tail parameters.

Third, the reformulation is *portable* across the symmetry-set machinery. Bloom's propagation arguments repeatedly compare convolutions and difference convolutions involving sets such as

$$S_\delta(A) = \{x \in G : (1_A \circ 1_A)(x) \geq \delta|A|\},$$

and one is led to estimates for quantities like $\|1_A \circ 1_S\|_{3/2}$ (or inner products involving such terms). In these steps, intermediate functions naturally appear at many scales, and the Lorentz framework isolates the only place where scale summation occurs: the passage from dyadic layers to a norm. Thus the same chain of inequalities used by Bloom can be run with 1_B replaced by general weights f, and with truncation arguments replaced by the identity defining $\|f\|_{3/2,1}$. The resulting statements are cleaner (no large-power tail losses), and they admit weighted analogues at essentially no additional cost.

In summary, by treating control as the operator norm of T_A from $\ell^{3/2,1}$ to ℓ^3, we obtain a single invariant $\mathsf{C}(A)$ that simultaneously encodes Bloom's third-moment bounds for indicators and provides the sharp extension needed for weighted and multi-scale arguments. The subsequent sections record the discrete Lorentz preliminaries and the basic inequalities that let us run this program in a self-contained way.

2 Preliminaries: convolutions, rearrangements, and discrete Lorentz spaces

Throughout we work on a fixed abelian group G, written additively. All functions $f, g, w\colon G \to [0, \infty)$ that appear in the sequel are assumed finitely supported unless explicitly stated otherwise; in particular, all sums over G are finite and rearrangements of summation are justified without further comment. For a finite set $A \subset G$ we write 1_A for its indicator function, and we use the discrete inner product

$$\langle f, g \rangle \ := \ \sum_{x \in G} f(x)g(x).$$

Convolutions on G

We use two closely related convolutions. The (additive) convolution is

$$(f * g)(x) \ := \ \sum_{y \in G} f(x-y)g(y),$$

and the (difference) convolution is

$$(f \circ g)(x) \ := \ \sum_{y \in G} f(x+y)g(y).$$

When $f = 1_A$ and $g = 1_B$ are indicators, $(1_A * 1_B)(x)$ counts representations $x = a + b$ with $a \in A$, $b \in B$, while $(1_A \circ 1_A)(x)$ counts representations $x = a - a'$ with $a, a' \in A$. We frequently use the involution $\tilde{f}(x) := f(-x)$, in terms of which one has the identities

$$f \circ g \ = \ f * \tilde{g}, \qquad \widetilde{f * g} \ = \ \tilde{g} * \tilde{f}.$$

These allow us to pass freely between $*$ and $\circ$ at the cost of a reflection.

Two elementary algebraic facts will be used repeatedly. First, convolution is associative and commutative (since G is abelian), and it is bilinear on finitely supported functions. Second, convolution interacts with the inner product via the standard adjointness relation: for finitely supported f, g, h,

$$\langle f * g, h \rangle \ = \ \langle g, \tilde{f} \circ h \rangle \ = \ \langle f, h \circ \tilde{g} \rangle.$$

We will also use the pointwise bounds

$$0 \le (f * g)(x) \le \|f\|_{\ell^1} \|g\|_{\ell^\infty}, \qquad \|f * g\|_{\ell^1} = \|f\|_{\ell^1} \|g\|_{\ell^1},$$

which are immediate from the definitions and nonnegativity.

Discrete ℓ^p spaces

For $p \in [1, \infty)$ we write

$$\|f\|_{\ell^p(G)} \ := \ \Big(\sum_{x \in G} |f(x)|^p \Big)^{1/p}, \qquad \|f\|_{\ell^\infty(G)} \ := \ \sup_{x \in G} |f(x)|.$$

The triangle inequality (Minkowski) and Hölder are used in their usual forms; for instance, if $1/p + 1/q = 1/r$ with $p, q, r \in [1, \infty]$, then

$$\|f * g\|_{\ell^r} \ \le \ \|f\|_{\ell^p} \|g\|_{\ell^q}$$

(Young's inequality), and if $1/p + 1/p' = 1$ then

$$\langle f, g \rangle \ \le \ \|f\|_{\ell^p} \|g\|_{\ell^{p'}}.$$

In practice we will apply Minkowski in ℓ^3 to sums of dyadic pieces, and we will apply Hölder to inner products arising from symmetry-set manipulations.

Decreasing rearrangements and distribution functions

To formulate the correct endpoint space for dyadic decompositions we recall the discrete decreasing rearrangement. If f is finitely supported, let $f^*(1) \geq f^*(2) \geq \cdots$ denote the nonincreasing rearrangement of $|f|$ on its support (extended by 0 thereafter). Equivalently, if we set the distribution function

$$\mu_f(t) \ := \ |\{x \in G : \ |f(x)| > t\}|,$$

then $f^*(n) = \inf\{t \geq 0 : \ \mu_f(t) < n\}$. Rearrangement invariance of the Lorentz norms will allow us to estimate quantities depending only on level sets of f, which is precisely what emerges from layer-cake expansions.

A convenient way to pass between a function and its dyadic level sets is to consider

$$E_k \ := \ \{x \in G : \ 2^k \leq f(x) < 2^{k+1}\}, \qquad k \in \mathbb{Z}.$$

For nonnegative finitely supported f we then have the pointwise comparison

$$f(x) \ \asymp \ \sum_{k \in \mathbb{Z}} 2^k 1_{E_k}(x),$$

with absolute implicit constants (indeed, $f \leq \sum_k 2^{k+1} 1_{E_k}$ and $f \geq \sum_k 2^k 1_{E_k}$). This elementary discretization is the starting point for all "no-cancellation" extensions from indicators to general weights.

Discrete Lorentz spaces $\ell^{p,1}(G)$

Let $p \in (1, \infty)$. The discrete Lorentz quasi-norm $\|f\|_{\ell^{p,1}(G)}$ is defined by

$$\|f\|_{\ell^{p,1}} \ := \ \sum_{n \geq 1} n^{1/p-1} f^*(n),$$

which is finite for finitely supported f. We will also use the weak Lorentz space $\ell^{p,\infty}(G)$ with quasi-norm

$$\|f\|_{\ell^{p,\infty}} \ := \ \sup_{n \geq 1} n^{1/p} f^*(n) \ = \ \sup_{t > 0} t \, \mu_f(t)^{1/p}.$$

Both $\ell^{p,1}$ and $\ell^{p,\infty}$ are rearrangement-invariant, and they sit at the endpoints of the scale $\ell^{p,1} \subset \ell^p \subset \ell^{p,\infty}$, with continuous embeddings and absolute constants. In particular, for finitely supported f,

$$\|f\|_{\ell^p} \ \leq \ \|f\|_{\ell^{p,1}}.$$

We emphasize that $\ell^{p,1}$ is the appropriate domain for dyadic layer-cake arguments: the quantity naturally produced by summing contributions from dyadic pieces agrees, up to absolute constants, with $\|f\|_{p,1}$.

Concretely, for $p = 3/2$ and nonnegative finitely supported f, the dyadic decomposition above yields the equivalence

$$\|f\|_{\ell^{3/2,1}} \asymp \sum_{k\in\mathbb{Z}} 2^k\,|E_k|^{2/3},$$

again with absolute implicit constants (the sum is finite when f is finitely supported and bounded). We will invoke this equivalence as the mechanism that replaces Bloom-style truncations: whenever we apply Minkowski to $\|1_A * f\|_{\ell^3}$ after decomposing f into $\sum_k 2^k 1_{E_k}$, the coefficient $\sum_k 2^k|E_k|^{2/3}$ is exactly the $\ell^{3/2,1}$ size of f.

Two simple computations will be used repeatedly. First, for a finite set $B \subset G$, the rearrangement of 1_B is $1_B^*(n) = 1$ for $1 \le n \le |B|$ and 0 thereafter, hence

$$\|1_B\|_{\ell^{3/2,1}} = \sum_{n=1}^{|B|} n^{-1/3} \asymp |B|^{2/3}.$$

Second, the quasi-triangle inequality for $\ell^{p,1}$ gives (for nonnegative finitely supported functions) a bound of the form

$$\Big\|\sum_j f_j\Big\|_{\ell^{p,1}} \le C_p \sum_j \|f_j\|_{\ell^{p,1}},$$

with C_p depending only on p; we will not need sharp dependence, only that the constants are absolute once p is fixed.

Lorentz–Hölder inequalities

Finally, we record the Lorentz refinements of Hölder that allow us to pair a strong Lorentz function with a weak Lorentz function. If $1 < p < \infty$ and p' is the conjugate exponent, then for finitely supported f, g we have

$$\langle f, g\rangle \le C_p\,\|f\|_{\ell^{p,1}}\,\|g\|_{\ell^{p',\infty}},$$

where C_p is an absolute constant depending only on p. This inequality is the discrete form of the standard Lorentz-space duality $(\ell^{p,1})^* = \ell^{p',\infty}$ (up to constants), and it is exactly what is needed when one factor is naturally controlled by level-set estimates.

We will also implicitly use that weak-ℓ^p control follows from level-set bounds: if $\mu_g(t) \le M^p t^{-p}$ for all $t > 0$, then $\|g\|_{\ell^{p,\infty}} \le M$. In the symmetry-set context, such estimates arise when one controls the size of $\{x : (1_A \circ 1_A)(x) \ge t\}$ at various thresholds t, and then pairs the resulting weak bounds with an $\ell^{p,1}$ quantity produced by a weighted convolution.

The preceding definitions and inequalities are the only analytic input we require. In the next section we apply them to the convolution operator

$T_A(f) = 1_A * f$, isolating the precise operator norm that is equivalent to Bloom's L^3-control and is stable under the dyadic manipulations described above.

3 The Lorentz control constant $\mathsf{C}(A)$

Fix a finite set $A \subset G$. We view convolution by 1_A as a linear operator

$$T_A \colon f \mapsto 1_A * f,$$

and we package its relevant endpoint mapping properties by the normalized Lorentz operator norm

$$\mathsf{C}(A) \ := \ \sup_{f \geq 0,\ f \not\equiv 0} \frac{\|1_A * f\|_{\ell^3(G)}}{|A|^{2/3}\, \|f\|_{\ell^{3/2,1}(G)}}.$$

The normalization by $|A|^{2/3} = \|1_A\|_{\ell^{3/2}}$ is chosen so that the trivial Young bound becomes scale-free, and so that $\mathsf{C}(A)$ coincides (in the next section) with Bloom's L^3-control parameter to the $1/3$ power. Since we restrict to $f \geq 0$, $\mathsf{C}(A)$ is tailored to applications in which cancellation plays no role; signed variants follow by inserting absolute values at the cost of harmless constants.

Basic bounds and invariances

We first note that $\mathsf{C}(A)$ is always finite and in fact bounded by an absolute constant. Indeed, Young's inequality with exponents $(3/2, 3/2, 3)$ gives

$$\|1_A * f\|_{\ell^3} \ \leq \ \|1_A\|_{\ell^{3/2}} \|f\|_{\ell^{3/2}} \ = \ |A|^{2/3}\, \|f\|_{\ell^{3/2}},$$

and the continuous embedding $\ell^{3/2,1} \subset \ell^{3/2}$ (equivalently, $\|f\|_{\ell^{3/2}} \leq \|f\|_{\ell^{3/2,1}}$) yields

$$\frac{\|1_A * f\|_{\ell^3}}{|A|^{2/3} \|f\|_{\ell^{3/2,1}}} \ \leq \ 1.$$

Thus

$$0 \ \leq \ \mathsf{C}(A) \ \leq \ 1. \tag{1}$$

A complementary lower bound (useful only as a sanity check) is obtained by testing on a point mass: if δ_0 is the indicator of $\{0\}$, then $1_A * \delta_0 = 1_A$ and $\|\delta_0\|_{\ell^{3/2,1}} = 1$, so

$$\mathsf{C}(A) \ \geq \ \frac{\|1_A\|_{\ell^3}}{|A|^{2/3}} \ = \ |A|^{-1/3}. \tag{2}$$

The quantity $\mathsf{C}(A)$ is invariant under the obvious symmetries of the group. If $x \in G$ and $A + x := \{a + x : a \in A\}$, then $1_{A+x} * f$ is a translate of $1_A * f$, hence has the same ℓ^3-norm; also $|A + x| = |A|$. Therefore

$$\mathsf{C}(A + x) = \mathsf{C}(A). \tag{3}$$

Similarly, if $-A := \{-a : a \in A\}$, then $1_{-A} * f = \widetilde{1_A} * f$, and reflection preserves ℓ^p-norms, so $\mathsf{C}(-A) = \mathsf{C}(A)$. More generally, any group automorphism ϕ preserves C provided we identify f with $f \circ \phi^{-1}$; we will only use the translation invariance (3).

Monotonicity in A

The operator T_A is monotone with respect to set inclusion in the sense appropriate to our normalization.

Lemma 3.1 (Monotonicity)**.** *If $A \subset A'$ are finite then*

$$\mathsf{C}(A) \ \le\ \Big(\frac{|A'|}{|A|}\Big)^{2/3} \mathsf{C}(A').$$

In particular, if $|A| = |A'|$ and $A \subset A'$ then $\mathsf{C}(A) \le \mathsf{C}(A')$.

Proof. For any $f \ge 0$ we have pointwise $1_A * f \le 1_{A'} * f$, hence $\|1_A * f\|_{\ell^3} \le \|1_{A'} * f\|_{\ell^3}$. Dividing by $|A|^{2/3}\|f\|_{\ell^{3/2,1}}$ gives

$$\frac{\|1_A * f\|_{\ell^3}}{|A|^{2/3}\|f\|_{\ell^{3/2,1}}} \ \le\ \Big(\frac{|A'|}{|A|}\Big)^{2/3} \frac{\|1_{A'} * f\|_{\ell^3}}{|A'|^{2/3}\|f\|_{\ell^{3/2,1}}}.$$

Taking the supremum over f yields the claim. □

We emphasize that monotonicity without the factor $(|A'|/|A|)^{2/3}$ is not the natural statement here, since $|A|$ is built into the normalization of $\mathsf{C}(A)$.

Subadditivity under disjoint unions

A key structural property, mirroring Bloom's decomposition steps, is subadditivity of C under disjoint unions.

Lemma 3.2 (Disjoint-union subadditivity)**.** *If $A = \bigsqcup_{i=1}^t A_i$ is a disjoint union of finite sets, then*

$$\mathsf{C}(A) \ \le\ \sum_{i=1}^t \mathsf{C}(A_i).$$

Proof. For $f \ge 0$ we have the decomposition $1_A * f = \sum_{i=1}^t (1_{A_i} * f)$. By Minkowski in ℓ^3,

$$\|1_A * f\|_{\ell^3} \ \le\ \sum_{i=1}^t \|1_{A_i} * f\|_{\ell^3} \ \le\ \sum_{i=1}^t \mathsf{C}(A_i)\,|A_i|^{2/3}\,\|f\|_{\ell^{3/2,1}}.$$

Since $|A_i| \le |A|$ for each i, we have $|A_i|^{2/3} \le |A|^{2/3}$, hence

$$\|1_A * f\|_{\ell^3} \ \le\ |A|^{2/3}\Big(\sum_{i=1}^t \mathsf{C}(A_i)\Big)\|f\|_{\ell^{3/2,1}}.$$

Dividing and taking the supremum over f gives the result. □

In later arguments, $\mathsf{C}(A)$ will therefore behave well under iterative pruning and decomposition: splitting off structured or sparse pieces increases C by at most an additive error, rather than forcing the introduction of auxiliary cutoffs.

Evaluation on indicators and the indicator-testing principle

The definition of $\mathsf{C}(A)$ is designed so that (i) it controls all indicator convolutions $1_A * 1_B$ uniformly in B, and (ii) conversely, testing only on indicators already captures $\mathsf{C}(A)$ up to absolute constants. The first direction is immediate: for any finite $B \subset G$,

$$\|1_A * 1_B\|_{\ell^3} \ \le\ \mathsf{C}(A)\,|A|^{2/3}\,\|1_B\|_{\ell^{3/2,1}}. \tag{4}$$

By the explicit computation of the Lorentz norm of an indicator,

$$\|1_B\|_{\ell^{3/2,1}} \ =\ \sum_{n=1}^{|B|} n^{-1/3} \ \asymp\ |B|^{2/3},$$

we deduce the convenient form

$$\|1_A * 1_B\|_{\ell^3} \ \ll\ \mathsf{C}(A)\,|A|^{2/3}\,|B|^{2/3}, \tag{5}$$

with an absolute implied constant depending only on the normalization of $\|\cdot\|_{\ell^{3/2,1}}$.

For the converse, we record the corresponding lower bound obtained by restricting the supremum in the definition of $\mathsf{C}(A)$ to indicators.

Lemma 3.3 (Indicator testing)**.** *There is an absolute constant $c > 0$ such that for every finite $A \subset G$,*

$$\mathsf{C}(A) \ \ge\ c \cdot \sup_{\varnothing \neq B \subset G \text{ finite}} \frac{\|1_A * 1_B\|_{\ell^3}}{|A|^{2/3}\,|B|^{2/3}}.$$

Proof. Fix $B \neq \varnothing$ and set $f = 1_B$ in the definition of $\mathsf{C}(A)$. Then

$$\mathsf{C}(A) \ \ge\ \frac{\|1_A * 1_B\|_{\ell^3}}{|A|^{2/3}\,\|1_B\|_{\ell^{3/2,1}}}.$$

Using $\|1_B\|_{\ell^{3/2,1}} \ll |B|^{2/3}$ yields the claim. □

Thus $\mathsf{C}(A)$ simultaneously governs the weighted regime (general $f \ge 0$) and contains, as a special case, the classical unweighted regime (indicators 1_B). The substantive point, proved in the next section, is that the upper bound (4) is not merely a consequence of the definition: it is exactly equivalent to Bloom's L^3-control inequality when the latter holds uniformly over all indicators 1_B. Put differently, passing from 1_B to general f incurs no truncation loss once one works in the correct domain space $\ell^{3/2,1}(G)$.

4 4. Equivalence theorem $\mathsf{C}(A) = \kappa(A)^{1/3}$: (a) indicator control implies Lorentz bound; (b) Lorentz bound implies indicator control; discussion of when polylog losses appear (only when converting $\ell^{3/2,1} \leftrightarrow \ell^{3/2}$).

Equivalence with Bloom's L^3-control parameter

We now identify the operator norm $\mathsf{C}(A)$ with Bloom's control parameter. Recall that $\kappa(A)$ is the minimal $\kappa \in (0,1]$ such that for every finite $B \subset G$ one has

$$\sum_{x\in G} \big(1_A * 1_B(x)\big)^3 \;\le\; \kappa\,|A|^2\,|B|^2, \tag{6}$$

or equivalently $\|1_A * 1_B\|_{\ell^3} \le \kappa^{1/3}|A|^{2/3}|B|^{2/3}$.

Theorem 4.1 (Equivalence of $\mathsf{C}(A)$ and $\kappa(A)^{1/3}$). *For every finite $A \subset G$ we have*

$$\mathsf{C}(A) \;=\; \kappa(A)^{1/3},$$

up to an absolute multiplicative constant depending only on the normalization of the Lorentz quasi-norm $\|\cdot\|_{\ell^{3/2,1}}$. In particular, after fixing the convention for $\|\cdot\|_{\ell^{3/2,1}}$, the two quantities determine one another by $\kappa(A) \asymp \mathsf{C}(A)^3$.

Proof. We prove the two implications separately.

(a) Indicator control $\Rightarrow$ Lorentz bound. Assume A satisfies (6) with constant κ. We claim that for every finitely supported $f \ge 0$,

$$\|1_A * f\|_{\ell^3} \;\ll\; \kappa^{1/3}\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}}. \tag{7}$$

Let $(E_k)_{k\in\mathbb{Z}}$ be the disjoint dyadic level sets

$$E_k \;:=\; \{x \in G : 2^k \le f(x) < 2^{k+1}\}.$$

Since f is finitely supported, only finitely many E_k are nonempty. Pointwise we have

$$f \;\le\; \sum_k 2^{k+1} 1_{E_k},$$

hence by positivity and linearity of convolution,

$$1_A * f \;\le\; \sum_k 2^{k+1}\,(1_A * 1_{E_k}).$$

Applying Minkowski's inequality in ℓ^3 gives

$$\|1_A * f\|_{\ell^3} \;\le\; \sum_k 2^{k+1}\,\|1_A * 1_{E_k}\|_{\ell^3}. \tag{8}$$

For each k, the hypothesis (6) with $B = E_k$ yields

$$\|1_A * 1_{E_k}\|_{\ell^3}^3 \;=\; \sum_x (1_A * 1_{E_k}(x))^3 \;\le\; \kappa\, |A|^2\, |E_k|^2,$$

and therefore

$$\|1_A * 1_{E_k}\|_{\ell^3} \;\le\; \kappa^{1/3}\, |A|^{2/3}\, |E_k|^{2/3}. \tag{9}$$

Substituting (9) into (8) gives

$$\|1_A * f\|_{\ell^3} \;\le\; 2\,\kappa^{1/3}\, |A|^{2/3} \sum_k 2^k\, |E_k|^{2/3}. \tag{10}$$

Finally, the Lorentz layer-cake identity for $\ell^{3/2,1}$ (in the discrete setting) asserts that

$$\|f\|_{\ell^{3/2,1}} \;\asymp\; \sum_k 2^k\, |E_k|^{2/3}, \tag{11}$$

with absolute implied constants. Combining (10) and (11) yields (7). Taking the supremum over $f \ge 0$ in the definition of $\mathsf{C}(A)$ gives $\mathsf{C}(A) \ll \kappa^{1/3}$.

(b) Lorentz bound $\Rightarrow$ indicator control. Conversely, assume $\mathsf{C}(A) < \infty$. For any finite $B \subset G$ we may test the defining inequality with $f = 1_B$, obtaining

$$\|1_A * 1_B\|_{\ell^3} \;\le\; \mathsf{C}(A)\, |A|^{2/3}\, \|1_B\|_{\ell^{3/2,1}}.$$

Using the explicit evaluation $\|1_B\|_{\ell^{3/2,1}} \asymp |B|^{2/3}$ (again with an absolute constant depending only on normalization), we deduce

$$\|1_A * 1_B\|_{\ell^3} \;\ll\; \mathsf{C}(A)\, |A|^{2/3}\, |B|^{2/3}.$$

Cubing both sides yields

$$\sum_x (1_A * 1_B(x))^3 \;=\; \|1_A * 1_B\|_{\ell^3}^3 \;\ll\; \mathsf{C}(A)^3\, |A|^2\, |B|^2,$$

so $\kappa(A) \ll \mathsf{C}(A)^3$, equivalently $\kappa(A)^{1/3} \ll \mathsf{C}(A)$.

Putting (a) and (b) together gives $\mathsf{C}(A) \asymp \kappa(A)^{1/3}$. If one fixes the normalization of $\|\cdot\|_{\ell^{3/2,1}}$ so that $\|1_B\|_{\ell^{3/2,1}} = |B|^{2/3}$ holds (up to the harmless endpoint convention for the rearrangement sum), then the implicit constants in the above comparison may be taken to be 1, yielding the stated identification $\mathsf{C}(A) = \kappa(A)^{1/3}$. $\square$

On the appearance of polylogarithmic losses

We emphasize that the equivalence in Theorem 4.1 is *tail-free* at the level of $\ell^{3/2,1}$: no truncation of small values of f is needed in the passage from indicators to general weights. Any polylogarithmic losses arise only when one insists on expressing results in terms of $\|f\|_{\ell^{3/2}}$ rather than $\|f\|_{\ell^{3/2,1}}$.

Indeed, from the embedding $\ell^{3/2,1} \subset \ell^{3/2}$ we always have $\|f\|_{\ell^{3/2}} \le \|f\|_{\ell^{3/2,1}}$, and therefore the Lorentz control bound implies the (formally weaker) estimate

$$\|1_A * f\|_{\ell^3} \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}} \;\ge\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{\ell^{3/2}}.$$

However, reversing this comparison—bounding $\|f\|_{\ell^{3/2,1}}$ in terms of $\|f\|_{\ell^{3/2}}$—necessarily depends on how many dyadic scales occur in f. Concretely, if $f = \sum_k 2^k 1_{E_k}$ with disjoint E_k, then by Hölder,

$$\sum_k 2^k |E_k|^{2/3} \;\le\; \Big(\sum_k (2^k)^{3/2} |E_k|\Big)^{2/3} \Big(\#\{k : E_k \neq \varnothing\}\Big)^{1/3} \;=\; \|f\|_{\ell^{3/2}} \Big(\#\text{scales}\Big)^{1/3}.$$

Thus whenever f is supported on at most m dyadic scales (for instance, when f is bounded between 1 and 2^m), we have $\|f\|_{\ell^{3/2,1}} \ll m^{1/3} \|f\|_{\ell^{3/2}}$. In applications, m is typically comparable to $\log(1/\eta)$ for an auxiliary cutoff η used to ignore very small values of f; the point is that Lorentz control removes the need to introduce such η in the first place, and therefore avoids the large tail terms that otherwise dominate when $\kappa(A)$ is small.

5 5. Tail-free function extension principle: replace Bloom's Lemma 2 with a Lorentz-strong estimate; derive clean corollaries for dyadic layer sets without κ^{100} terms.

Tail-free function extension principle

A recurring technical step in Bloom's arguments is to pass from an ℓ^3 estimate for convolutions with indicators 1_B to an estimate for convolutions with general nonnegative weights f. In Bloom's formulation this passage is implemented by a truncation-and-pigeonhole device (his Lemma 2), which introduces an auxiliary lower cutoff to discard small values of f; the discarded tail is then bounded crudely and re-enters later as a large power of κ^{-1} (of the schematic form κ^{100} in the denominator). For our purposes it is preferable to replace this mechanism by a statement which is (i) linear in f, (ii) stable under arbitrary superpositions of dyadic layers, and (iii) exact up to absolute constants. The Lorentz formulation of control provides precisely this.

Proposition 5.1 (Tail-free extension to weights)**.** *For every finite $A \subset G$ and every finitely supported $f \ge 0$ we have*

$$\|1_A * f\|_{\ell^3(G)} \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}(G)}. \tag{12}$$

Equivalently, if A satisfies Bloom control (6) *with constant κ, then*

$$\|1_A * f\|_{\ell^3(G)} \;\ll\; \kappa^{1/3}\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}(G)}. \tag{13}$$

This is immediate from the definition of $\mathsf{C}(A)$, and (13) is the same statement after invoking Theorem 4.1 to identify $\mathsf{C}(A) \asymp \kappa^{1/3}$ (with normalization-dependent absolute constants). We isolate (12) as a "principle" because it is the exact replacement for every occurrence of Bloom's function-extension lemma: whenever an argument uses κ-control only through bounds of the form $\|1_A * 1_B\|_3$ and then extends from 1_B to a weight, one may instead work directly with $\ell^{3/2,1}$ and insert (12) without any cutoff.

The point is not merely aesthetic: the Lorentz quasi-norm is the natural bookkeeping device for dyadic layer-cake decompositions, and so the extension (12) is compatible with arbitrary superpositions of scales, with no remainder terms. We record this explicitly.

Corollary 5.2 (Dyadic superposition bound). *Let $(E_k)_{k\in\mathbb{Z}}$ be pairwise disjoint finite subsets of G, and let $(\lambda_k)_{k\in\mathbb{Z}}$ be nonnegative coefficients with only finitely many nonzero terms. Set*

$$f \;:=\; \sum_k \lambda_k\, 1_{E_k}.$$

Then

$$\|1_A * f\|_{\ell^3} \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}} \;\ll\; \mathsf{C}(A)\,|A|^{2/3}\sum_k \lambda_k\,|E_k|^{2/3}. \tag{14}$$

In particular, for the dyadic choice $\lambda_k = 2^k$ and $E_k = \{x : 2^k \le f(x) < 2^{k+1}\}$ one has

$$\|1_A * f\|_{\ell^3} \;\ll\; \mathsf{C}(A)\,|A|^{2/3}\sum_k 2^k\,|E_k|^{2/3} \;\asymp\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{\ell^{3/2,1}}. \tag{15}$$

The second inequality in (14) is simply the Lorentz estimate $\|f\|_{3/2,1} \ll \sum_k \lambda_k |E_k|^{2/3}$, which is a direct consequence of the definition of $\|\cdot\|_{3/2,1}$ by decreasing rearrangement (or, equivalently, of the layer-cake identity (11) when λ_k are dyadic). The crucial feature of (14) is that the contribution of each layer is additive, and there is no need to locate a "dominant scale" by pigeonholing, nor to throw away the smallest layers by truncation.

A convenient way to phrase the elimination of tail errors is to note that $\ell^{3/2,1}$ is designed so that truncations are controlled monotonically: if $0 \le f_1 \le f_2$ pointwise, then $\|f_1\|_{3/2,1} \le \|f_2\|_{3/2,1}$. Consequently, every truncated version of Bloom's lemma that bounds $\|1_A * (f\,1_{\{f\ge\eta\}})\|_3$ for some cutoff $\eta > 0$ is subsumed by (12), without any need to track how the complementary part $f\,1_{\{f<\eta\}}$ is treated. For example, for any $\eta > 0$,

$$\|1_A * (f\,1_{\{f\ge\eta\}})\|_3 \;\le\; \|1_A * f\|_3 \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{3/2,1}, \tag{16}$$

and similarly

$$\|1_A * (f\,1_{\{f<\eta\}})\|_3 \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\,1_{\{f<\eta\}}\|_{3/2,1} \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{3/2,1}. \tag{17}$$

Thus the small values of f are never problematic at the level of $\ell^{3/2,1}$; they are accounted for exactly in the norm. This is the precise sense in which the Lorentz formulation is tail-free.

We also note that (12) is insensitive to replacing additive convolution by difference convolution, up to the harmless reflection $f^{\sim}(x) := f(-x)$. Indeed, for any $g \geq 0$,

$$1_A \circ g \ = \ 1_A * g^{\sim},$$

and $\|g^{\sim}\|_{3/2,1} = \|g\|_{3/2,1}$ by rearrangement invariance. Hence

$$\|1_A \circ g\|_3 \ \leq \ \mathsf{C}(A)\,|A|^{2/3}\,\|g\|_{3/2,1}. \tag{18}$$

We will use (18) systematically when estimating weighted symmetry expressions.

Finally, we emphasize how (12) interfaces with later arguments. In propagation and symmetry-set manipulations one repeatedly encounters weights obtained by multiplying or summing indicators across several scales (for instance, functions like $w = \sum_i 1_{B_i}$ or $w = \sum_k 2^k 1_{E_k}$ arising from dyadic decompositions of representation functions). The estimate (14) shows that, as long as $\|w\|_{3/2,1}$ is controlled, the ℓ^3 norm of $1_A * w$ is controlled with no additional bookkeeping. In particular, there is no analogue of a κ^{100}-type penalty attached to the number of dyadic scales present in w; any dependence on scale complexity enters only if one insists on expressing $\|w\|_{3/2,1}$ in terms of $\|w\|_{3/2}$, as discussed previously.

We will treat (12)–(18) as the basic replacement rule for Bloom's function-extension lemma. In the next section we apply this replacement to the symmetry-set machinery, where the relevant weights are naturally produced by thresholding and dyadic decompositions of convolutions, and where the absence of tail terms materially simplifies the dependence on auxiliary cutoffs.

6 6. Weighted symmetry-set machinery: define weighted symmetry sets via thresholds on $1_A \circ w$; prove Lorentz versions of the key bounds on $\|1_A \circ 1_S\|_{3/2}$ and the auxiliary inner-product inequalities.

Weighted symmetry-set machinery

In the symmetry-set portion of Bloom's argument one isolates a set of shifts on which a difference convolution is *popular*, and then exploits this popularity through an inner-product identity and an ℓ^3 bound for a secondary convolution. The Lorentz formulation lets us run the same mechanism with general weights, and we record the basic templates in a form we will use later.

Let $w : G \to [0, \infty)$ be finitely supported. We write

$$u \ := \ 1_A \circ w, \qquad \text{so that} \qquad u(x) = \sum_{a \in A} w(x+a).$$

For a parameter $\delta \in (0, 1]$ we define the *weighted symmetry set at level* δ by

$$S_\delta(A; w) \ := \ \big\{x \in G : \ (1_A \circ w)(x) \ \geq \ \delta \, \|w\|_{\ell^1(G)}\big\}. \tag{19}$$

When $w = 1_A$ this recovers the usual symmetry set $S_\delta(A)$, up to the harmless normalisation $\|1_A\|_1 = |A|$.

The starting point is the adjointness of difference convolution under the ℓ^2 pairing. We will use it in the form below.

Lemma 6.1 (Adjointness identity)**.** *For finitely supported* $f, g, h \geq 0$ *one has*

$$\langle 1_A \circ f, \, g\rangle \ = \ \langle 1_A \circ g, \, f\rangle.$$

In particular, for $S \subset G$ *finite and* $w \geq 0$*,*

$$\langle 1_A \circ w, \, 1_S\rangle \ = \ \langle 1_A \circ 1_S, \, w\rangle. \tag{20}$$

Proof. Expanding and changing order of summation,

$$\langle 1_A \circ f, g\rangle = \sum_x \sum_y 1_A(x+y) f(y) g(x) = \sum_y f(y) \sum_x 1_A(x+y) g(x) = \langle 1_A \circ g, f\rangle.$$

The specialisation (20) is immediate. □

Popularity on $S_\delta(A; w)$ converts directly into a lower bound for this inner product.

Lemma 6.2 (Popularity lower bound)**.** *Let* $S = S_\delta(A; w)$ *as in* (19)*. Then*

$$\langle 1_A \circ w, \, 1_S\rangle \ \geq \ \delta \, \|w\|_1 \, |S|. \tag{21}$$

Proof. By definition of S, we have $(1_A \circ w)(x) \geq \delta \|w\|_1$ for all $x \in S$. Hence

$$\langle 1_A \circ w, 1_S\rangle = \sum_{x \in S} (1_A \circ w)(x) \ \geq \ \sum_{x \in S} \delta \|w\|_1 = \delta \|w\|_1 |S|.$$

□

The point of (20) is that it moves the set S into the *second* factor $1_A \circ 1_S$, which is exactly where the ℓ^3 control input applies. For the lower bound, we combine (20)–(21) with Hölder (or Lorentz–Hölder, if one wishes to work with weak spaces).

Proposition 6.3 ($\ell^{3/2}$ lower bound for $1_A \circ 1_S$)**.** *Let $w \geq 0$ be finitely supported and suppose $w \in \ell^3(G)$. Let $S = S_\delta(A; w)$. Then*

$$\|1_A \circ 1_S\|_{\ell^{3/2}(G)} \;\geq\; \frac{\delta \,\|w\|_1 \,|S|}{\|w\|_3}. \tag{22}$$

In particular, taking $w = 1_A$ gives

$$\|1_A \circ 1_{S_\delta(A)}\|_{3/2} \;\geq\; \delta \,|A|^{2/3} \,|S_\delta(A)|. \tag{23}$$

Proof. By Lemma 6.1 and Lemma 6.2,

$$\delta\|w\|_1|S| \leq \langle 1_A \circ w, 1_S\rangle = \langle 1_A \circ 1_S, w\rangle.$$

Applying Hölder with exponents $(3/2, 3)$ yields

$$\langle 1_A \circ 1_S, w\rangle \leq \|1_A \circ 1_S\|_{3/2} \,\|w\|_3,$$

which implies (22). For $w = 1_A$, we have $\|w\|_1 = |A|$ and $\|w\|_3 = |A|^{1/3}$, giving (23). □

A variant that is occasionally convenient is to replace $\|w\|_3$ by a weak norm, using the Lorentz Hölder template (Lemma 4 in the global context). Namely, from $\langle F, w\rangle \leq \|F\|_{3/2,1}\|w\|_{3,\infty}$ we obtain

$$\|1_A \circ 1_S\|_{3/2,1} \;\geq\; \frac{\delta \,\|w\|_1 \,|S|}{\|w\|_{3,\infty}},$$

which is useful when w is a superposition of indicators. We will not emphasise this further, but it fits seamlessly with the same bookkeeping.

We next record the complementary ℓ^3 bound for $1_A \circ 1_S$, which is the exact point where Lorentz control replaces Bloom's truncation-based extension lemma. This is simply the difference-convolution version (18) specialised to indicators.

Proposition 6.4 (ℓ^3 control for symmetry convolutions)**.** *For every finite $S \subset G$,*

$$\|1_A \circ 1_S\|_{\ell^3(G)} \;\leq\; \mathsf{C}(A)\,|A|^{2/3}\,\|1_S\|_{\ell^{3/2,1}(G)} \;\asymp\; \mathsf{C}(A)\,|A|^{2/3}\,|S|^{2/3}. \tag{24}$$

More generally, for any finitely supported $v \geq 0$,

$$\|1_A \circ v\|_3 \;\leq\; \mathsf{C}(A)\,|A|^{2/3}\,\|v\|_{3/2,1}. \tag{25}$$

The pair of estimates (22) and (24) is the basic symmetry-set input: the popularity condition forces $1_A \circ 1_S$ to be large in $\ell^{3/2}$, while Lorentz control bounds the same function in ℓ^3. The remaining steps in the symmetry-set machine consist of inserting these two bounds into whatever interpolation or

Cauchy–Schwarz estimate is relevant at the given stage. For instance, since $\|1_A \circ 1_S\|_1 = |A|\,|S|$, Cauchy–Schwarz gives

$$\|1_A \circ 1_S\|_{3/2}^{3/2} = \sum_x (1_A \circ 1_S)(x)\,(1_A \circ 1_S)(x)^{1/2} \le \|1_A \circ 1_S\|_1^{1/2}\,\|1_A \circ 1_S\|_2, \quad (26)$$

hence

$$\|1_A \circ 1_S\|_2 \;\ge\; \frac{\|1_A \circ 1_S\|_{3/2}^{3/2}}{(|A|\,|S|)^{1/2}}. \quad (27)$$

Combining (27) with (23) yields an explicit lower bound on $\|1_A \circ 1_{S_\delta(A)}\|_2$ in terms of δ, $|A|$, and $|S_\delta(A)|$; in the weighted setting one obtains the same conclusion with $|A|$ replaced by $\|w\|_1^2/\|w\|_3^2$ via (22). These are exactly the kinds of inner-product manipulations that occur repeatedly in Bloom's propagation steps, and they require no truncation once (25) is available.

Finally, we note that (19) is only one convenient normalisation. In applications we will sometimes define S by a threshold of the form $(1_A \circ w)(x) \ge \tau$ for an absolute τ (or for τ depending on other parameters). The same formal argument gives

$$\tau\,|S| \;\le\; \langle 1_A \circ w, 1_S\rangle = \langle 1_A \circ 1_S, w\rangle \le \|1_A \circ 1_S\|_{3/2}\,\|w\|_3,$$

and hence $\|1_A \circ 1_S\|_{3/2} \ge \tau|S|/\|w\|_3$, with the ℓ^3 upper bound for $1_A \circ 1_S$ still furnished by (24). Thus the symmetry-set machinery is stable under arbitrary weighted choices, provided we measure the weights in spaces compatible with the Lorentz control inequality.

7 7. Propagation results in weighted/asymmetric form: restate Bloom's main applications (energy, sum/difference sets, BSG extraction, convex-function decomposition) with weights/measures and Lorentz control.

Propagation results in weighted/asymmetric form

We now record a convenient way to package Bloom's propagation outputs so that (a) weights are permitted from the outset, and (b) the only input is Lorentz control, with no auxiliary truncation parameters. The guiding principle is that every time Bloom invokes his function-extension lemma (which replaces an indicator 1_B by a general f at the cost of a large tail term), we instead apply the sharp Lorentz bound

$$\|1_A * f\|_3 \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|f\|_{3/2,1},$$

and then proceed through the same symmetry-set and interpolation steps as in the unweighted argument.

For flexibility it is useful to allow a general nonnegative weight $a : G \to [0, \infty)$ in place of 1_A. Accordingly we define the (normalised) Lorentz control constant of a by

$$\mathsf{C}(a) \;:=\; \sup_{f \geq 0,\ f \not\equiv 0} \frac{\|a * f\|_3}{\|a\|_1^{2/3}\,\|f\|_{3/2,1}}. \tag{28}$$

When $a = 1_A$ this recovers $\mathsf{C}(A)$. Moreover, by testing on indicators and using Lemma 1, one obtains the weighted analogue of Bloom's cubic moment control:

$$\|a * 1_B\|_3^3 \;\ll\; \mathsf{C}(a)^3\,\|a\|_1^2\,|B|^2 \qquad \text{for all finite } B \subset G, \tag{29}$$

with the same normalisation constants as in the indicator case. Thus, wherever Bloom's arguments are expressed in terms of $\kappa(A)$, we may equivalently work with $\mathsf{C}(a)$ via the identity $\kappa(a) = \mathsf{C}(a)^3$ (up to the fixed normalisation constant coming from Lemma 1).

Energy and sum/difference bounds with weights. Bloom's main quantitative outputs bound the additive energy and force growth of sumsets and difference sets. In weighted form, it is most natural to phrase the conclusions in terms of the weighted energy

$$E(a) \;:=\; \|a * a\|_2^2 \;=\; \sum_x (a * a(x))^2,$$

and the associated "effective support sizes"

$$|a + a|_{\text{eff}} \;:=\; \frac{\|a\|_1^4}{\|a * a\|_2^2}, \qquad |a - a|_{\text{eff}} \;:=\; \frac{\|a\|_1^4}{\|a \circ a\|_2^2}, \tag{30}$$

which coincide with $|A+A|$ and $|A-A|$ when $a = 1_A$ (up to the usual Cauchy–Schwarz comparison between support size and L^2 mass). The propagation mechanism in Bloom's proof uses only: (i) popularity bounds of the type recorded in the previous subsection, (ii) Cauchy–Schwarz/Hölder/interpolation, and (iii) L^3-control for secondary convolutions. Since (iii) is now available in the exact Lorentz form for arbitrary weights, the entire chain of implications carries over verbatim.

In particular, substituting $\kappa = \mathsf{C}^3$ throughout Bloom's exponents and keeping track only of the homogeneities forced by the definitions, we obtain the following template.

Corollary 7.1 (Weighted form of Bloom's Theorem 1)**.** *Let $a \geq 0$ be finitely supported and let $\varepsilon > 0$. Then*

$$E(a) \;\ll_\varepsilon\; \mathsf{C}(a)^{81/50-\varepsilon}\,\|a\|_1^3\,\|a\|_\infty, \tag{31}$$

and consequently

$$|a+a|_{\mathrm{eff}} \;\gg_\varepsilon\; \mathsf{C}(a)^{-33/19+\varepsilon}\,\frac{\|a\|_1}{\|a\|_\infty}, \qquad |a-a|_{\mathrm{eff}} \;\gg_\varepsilon\; \mathsf{C}(a)^{-7518/4175+\varepsilon}\,\frac{\|a\|_1}{\|a\|_\infty}. \tag{32}$$

If $a = 1_A$ these reduce to Bloom's bounds stated in terms of $\mathsf{C}(A)$ *(equivalently $\kappa(A)$).*

The factor $\|a\|_\infty$ is the natural scaling correction: under dilation $a \mapsto ta$, one has $E(ta) = t^4 E(a)$ while $\mathsf{C}(ta) = t^{1/3}\mathsf{C}(a)$, and (31) is precisely homogeneous. In typical applications one has $\|a\|_\infty \le 1$ (weights are sub-indicators or probability measures up to scaling), and then $\|a\|_\infty$ may be suppressed.

Asymmetric consequences. The same propagation steps can be run asymmetrically once one notices that the only L^3 input needed is control of the form $\|a * f\|_3 \lesssim \mathsf{C}(a)\|a\|_1^{2/3}\|f\|_{3/2,1}$, and the remainder of the argument is bookkeeping with difference convolutions and symmetry sets (which are inherently asymmetric). For instance, for two weights $a, b \ge 0$ we may consider the mixed energy

$$E(a,b) \;:=\; \|a * b\|_2^2 \;=\; \sum_x (a * b(x))^2,$$

and the effective sumset size $\|a\|_1^2\|b\|_1^2/\|a * b\|_2^2$. Whenever Bloom's proof treats B as an auxiliary set and uses control of A only through $\|1_A * 1_B\|_3$, we may replace 1_B by an arbitrary b with $\|b\|_{3/2,1}$ bounded and obtain the same conclusions with $|B|^{2/3}$ replaced by $\|b\|_{3/2,1}$. Concretely, the input $\|1_A * 1_B\|_3 \ll \mathsf{C}(A)|A|^{2/3}|B|^{2/3}$ becomes

$$\|1_A * b\|_3 \;\le\; \mathsf{C}(A)\,|A|^{2/3}\,\|b\|_{3/2,1}, \tag{33}$$

and no other step distinguishes b from an indicator except for this norm conversion. Thus the asymmetric, weighted variants of all intermediate "propagation inequalities" in Bloom's Sections 3–6 hold with the same exponents and with $\ell^{3/2,1}$ in place of $\ell^{3/2}$.

BSG extraction from an L^3 certificate. Bloom's "BSG-from-L^3" mechanism (Theorem 15-style) starts from a certificate that some convolution has abnormally large L^3 mass and then produces a structured subset (or approximate group) on which additive energy concentrates. In our language, the hypothesis is naturally phrased as the existence of a weight $b \ge 0$ for which $\|1_A * b\|_3$ is large compared to the normalised scale $|A|^{2/3}\|b\|_{3/2,1}$. The Lorentz framework is particularly well suited here: if one writes b as a dyadic superposition, then $\|b\|_{3/2,1}$ captures exactly the relevant layer-cake sum, and the extraction proceeds without the need to ignore small values of b by a cutoff depending on κ. In other words, any conclusion in Bloom's extraction theorem whose proof previously depended on a truncation lemma may be re-run with $\|b\|_{3/2,1}$ replacing $\|b\|_{3/2}$, yielding the same exponent-level dependence on $\mathsf{C}(A)$ and eliminating the large-κ tail losses. If one insists on an $\ell^{3/2}$ hypothesis on b rather than $\ell^{3/2,1}$, then the only additional losses are those incurred when bounding $\|b\|_{3/2,1}$ by $\|b\|_{3/2}$, which are at worst polylogarithmic under mild bounded-range hypotheses on b.

Decompositions (convex-function style) and pruning. Finally, Bloom repeatedly decomposes a set into pieces and propagates control estimates across the decomposition. In the Lorentz-operator formulation this is expressed cleanly by subadditivity at the level of C: whenever $a = \sum_{i=1}^{t} a_i$ with $\mathrm{supp}(a_i)$ pairwise disjoint, we have

$$\mathsf{C}(a) \;\leq\; \sum_{i=1}^{t} \mathsf{C}(a_i),$$

which is the weighted analogue of Lemma 3. This is the appropriate interface for convex-function-type decompositions: one isolates components a_i that are individually "geometric" (hence have small $\mathsf{C}(a_i)$ via incidence bounds) and controls the remainder by a crude estimate, while keeping a linear bookkeeping of the control constants rather than suffering powers of κ. In particular, iterative pruning arguments that repeatedly discard a small portion of mass and re-run the propagation step are stable under weights, since each iteration only requires applying (28) to the current weight and using the symmetry-set templates already recorded.

The net effect is that Bloom's entire propagation apparatus admits a uniform weighted/asymmetric restatement with C as the controlling invariant and $\ell^{3/2,1}$ as the natural domain space. This does not, by itself, improve any exponent, but it makes the framework robust under weights and removes all dependence on ad hoc truncation parameters.

8 8. Examples and sharpness: convex sets (Szemerédi–Trotter) yield $\mathsf{C}(A) \lesssim |A|^{-1/3}$; unions (subadditivity); toy weighted examples where Lorentz beats $\ell^{3/2}$; discussion of limitations (cannot by itself improve sumset exponent).

Examples and sharpness

We record a few examples illustrating (i) how geometric information about A can be injected through upper bounds for $\mathsf{C}(A)$, (ii) how the operator viewpoint interacts cleanly with decompositions, and (iii) why $\ell^{3/2,1}$ is the correct domain space if one wants a genuinely tail-free extension from indicators to general weights.

Convex sets in $\mathbb{R}$ and Szemerédi–Trotter. Let $A \subset \mathbb{R}$ be a finite strictly convex set (e.g. $A = \{a_1 < \cdots < a_n\}$ with strictly increasing successive differences). A standard incidence argument (going back to Elekes–Nathanson–Ruzsa and refined in many later works) shows that additive representation counts for such A have strong L^3-control. Concretely, one obtains an inequality of the form

$$\sum_x (1_A * 1_B(x))^3 \ \ll\ |A|\,|B|^2 \qquad \text{for all finite } B \subset \mathbb{R}, \tag{34}$$

with an absolute implied constant. (One way to view (34) is that solutions to $a_1 + b_1 = a_2 + b_2 = a_3 + b_3$ with $a_i \in A$, $b_i \in B$, are controlled by incidences between the point set $B \times B$ and a family of lines determined by differences in A; strict convexity provides the non-degeneracy needed to apply Szemerédi–Trotter without large multiplicity losses.)

Comparing (34) with Bloom's normalisation $\sum_x (1_A * 1_B(x))^3 \le \kappa(A)\,|A|^2|B|^2$, we read off $\kappa(A) \ll |A|^{-1}$. By Theorem A this is equivalent to

$$\mathsf{C}(A) \ =\ \kappa(A)^{1/3} \ \ll\ |A|^{-1/3}. \tag{35}$$

Thus convexity supplies a small operator norm, and the Lorentz extension then immediately yields, for all finitely supported $f \ge 0$,

$$\|1_A * f\|_3 \ \ll\ |A|^{1/3}\,\|f\|_{3/2,1}.$$

We emphasise that (35) is the natural scaling: $\|1_A * f\|_3$ has the same homogeneity as $|A|^{2/3}\|f\|_{3/2,1}$, and the additional gain $|A|^{-1/3}$ is exactly what the incidence estimate provides. In particular, in any propagation argument whose quantitative output is a negative power of $\mathsf{C}(A)$, convexity immediately forces polynomial growth of sumsets/difference sets as a function of $|A|$.

Disjoint unions and subadditivity (sharpness of the linear bookkeeping). Suppose $A = \bigsqcup_{i=1}^{t} A_i$ is a disjoint union. Then $1_A * f = \sum_i 1_{A_i} * f$, and by the triangle inequality in ℓ^3 we obtain

$$\|1_A * f\|_3 \ \le\ \sum_{i=1}^{t} \|1_{A_i} * f\|_3 \ \le\ \Big(\sum_{i=1}^{t} \mathsf{C}(A_i)\Big)\,|A|^{2/3}\,\|f\|_{3/2,1},$$

after the harmless renormalisation $\|1_{A_i}\|_1^{2/3} \le \|1_A\|_1^{2/3}$. Taking the supremum over f yields

$$\mathsf{C}(A) \ \le\ \sum_{i=1}^{t} \mathsf{C}(A_i), \qquad \text{equivalently} \qquad \kappa(A)^{1/3} \ \le\ \sum_{i=1}^{t} \kappa(A_i)^{1/3}. \tag{36}$$

This inequality is frequently close to sharp. For instance, if the pieces A_i are widely separated in the ambient group so that the functions $1_{A_i} * f$

have essentially disjoint supports for the class of test functions f relevant to a given application, then $\|\sum_i 1_{A_i} * f\|_3$ behaves like $(\sum_i \|1_{A_i} * f\|_3^3)^{1/3}$, and the linear bound (36) is the correct robust statement one can guarantee without further structural input. In particular, this linear bookkeeping is precisely what one wants in iterative pruning/decomposition schemes: each step can be analysed at the level of C without introducing artificial powers of κ.

It is instructive to contrast (36) with what one would obtain by working only with $\kappa(A)$ directly: since $\kappa = \mathsf{C}^3$, a linear bound for C corresponds to a cubic-type bound for κ, and keeping the argument at the operator-norm level is what prevents such nonlinearities from accumulating uncontrollably across many decomposition steps.

A toy weighted example: why $\ell^{3/2,1}$ is the tail-free endpoint. The Lorentz norm $\|\cdot\|_{3/2,1}$ is genuinely stronger than $\|\cdot\|_{3/2}$, and this strength is not an artefact: it is exactly what eliminates tail bookkeeping. We spell this out with a simple dyadic construction.

Fix pairwise disjoint sets $B_1, \dots, B_m \subset G$ and a large parameter N, with

$$|B_j| \;\sim\; N\,2^{-3j/2} \qquad (1 \le j \le m),$$

so that the sizes decay geometrically. Define a weight

$$f \;:=\; \sum_{j=1}^{m} 2^j\, 1_{B_j}.$$

A direct computation gives

$$\|f\|_{3/2}^{3/2} \;=\; \sum_{j=1}^{m} (2^j)^{3/2} |B_j| \;\sim\; \sum_{j=1}^{m} 2^{3j/2} \cdot N\,2^{-3j/2} \;\sim\; mN,$$

and hence $\|f\|_{3/2} \sim (mN)^{2/3}$. On the other hand, Lemma 2 (or a one-line rearrangement computation) yields

$$\|f\|_{3/2,1} \;\asymp\; \sum_{j=1}^{m} 2^j\, |B_j|^{2/3} \;\sim\; \sum_{j=1}^{m} 2^j \cdot (N^{2/3} 2^{-j}) \;\sim\; m\,N^{2/3}.$$

Therefore

$$\frac{\|f\|_{3/2,1}}{\|f\|_{3/2}} \;\sim\; \frac{m\,N^{2/3}}{(mN)^{2/3}} \;=\; m^{1/3}. \tag{37}$$

Thus $\|f\|_{3/2}$ undercounts the dyadic complexity by a factor $m^{1/3}$. If one attempted to replace the Lorentz bound $\|1_A * f\|_3 \;\le\; \mathsf{C}(A)|A|^{2/3}\|f\|_{3/2,1}$ by an $\ell^{3/2}$-based inequality with the same constant $\mathsf{C}(A)$, then (37) shows that one would necessarily lose a factor $\gtrsim m^{1/3}$ on this class of examples unless one introduces additional hypotheses (bounded range, truncations, or a priori control of the number of active dyadic scales). This is exactly the phenomenon behind Bloom-style tail terms: the truncation parameter is, in effect, a proxy for the number of relevant scales, whereas $\ell^{3/2,1}$ records the scale sum intrinsically.

In applications where the weight f is known to be essentially bounded, say $0 \le f \le M$, one can of course trade Lorentz for Lebesgue at polylogarithmic cost (since only $\ll \log M$ dyadic levels occur). The point is that such losses come from external information about the range of f, not from the convolution estimate itself.

Limitations and what is (not) made sharper. The passage from $\kappa(A)$ to $\mathsf{C}(A)$ is an exact repackaging (Theorem A), and accordingly it cannot, by itself, improve any exponent in Bloom's energy or sumset bounds: any numerical gain would have to come from new input giving a smaller $\mathsf{C}(A)$ for the specific class of sets/weights under consideration (incidence geometry, Fourier-analytic bounds, structural information, or an inverse theorem for near-extremisers of C). What the Lorentz formulation does provide is stability: once an estimate for $\mathsf{C}(a)$ is available, it propagates through the entire machinery without additional losses from truncations, and it does so uniformly for weighted and asymmetric configurations. In this sense C is best viewed as an interface: it isolates the genuinely difficult geometric/combinatorial input (bounding an operator norm) from the downstream additive-combinatorial deductions.

9 9. Outlook (2026 directions): finite-field L^4 analogues; automated inequality search with Lorentz primitives; potential stability/inverse theory for Lorentz control.

Outlook (directions for 2026)

We conclude by indicating three directions where the Lorentz-operator formulation seems particularly well suited: (a) higher-moment analogues (notably an ℓ^4 theory over finite fields), (b) semi-automated proof search and optimisation once truncations are eliminated at the axiomatic level, and (c) a stability/inverse theory for near-extremisers of Lorentz control.

1. Finite-field ℓ^4 analogues and higher moments. The passage from Bloom's L^3-certificate to the operator norm $\mathsf{C}(A)$ suggests an obvious generalisation: for $q > 2$ one may define a q-moment control parameter

$$\kappa_q(A) \;:=\; \inf\Big\{\kappa > 0 : \; \sum_x (1_A * 1_B(x))^q \;\le\; \kappa\, |A|^{q-1}|B|^{q-1} \;\text{ for all finite } B\Big\},$$

and the corresponding Lorentz control constant

$$\mathsf{C}_q(A) \ := \sup_{f\geq 0,\ f\not\equiv 0} \frac{\|1_A * f\|_{\ell^q(G)}}{|A|^{1-1/q}\,\|f\|_{\ell^{q'},1(G)}}, \qquad q' = \frac{q}{q-1}.$$

At the level of formal manipulations, the proof of Theorem A is not special to $q = 3$: the same dyadic layer-cake argument gives $\mathsf{C}_q(A) = \kappa_q(A)^{1/q}$ (up to the normalisation constant coming from $\|1_B\|_{q',1}$). What changes, and where new input is needed, is in obtaining nontrivial upper bounds for $\kappa_q(A)$ (equivalently $\mathsf{C}_q(A)$) in concrete settings.

The case $q = 4$ over $G = \mathbb{F}_p^n$ looks especially promising. Several incidence bounds in finite fields (point–line and point–plane, in the spirit of Rudnev-type arguments) naturally control fourth moments of representation functions, and the absence of ordering makes truncation-based arguments particularly awkward; the endpoint Lorentz formulation avoids this from the outset. Moreover, ℓ^4 norms interface more directly with additive energy through identities such as

$$E(A) \ = \ \|1_A * 1_A\|_2^2, \qquad \|1_A * 1_B\|_4^4 \ = \ \sum_x (1_A * 1_B(x))^4,$$

and thus with higher-order energies and certain Gowers-type quantities. A systematic ℓ^4 control theory might therefore provide an alternate route to (or refinement of) energy increment arguments in characteristic p, where one often needs to handle weighted convolutions arising from density increments or random sampling.

Two concrete problems emerge. First, identify natural geometric hypotheses on $A \subset \mathbb{F}_p^n$ (e.g. algebraic curves/surfaces with suitable non-degeneracy, Cartesian products with expansion properties, or sets avoiding subfield structure) that force polynomial savings in $\mathsf{C}_4(A)$ relative to the trivial bound. Second, understand to what extent $\mathsf{C}_4(A)$ controls the propagation machinery that is currently organised around ℓ^3: which parts improve, which become weaker, and which require hybrid arguments (for instance, combining ℓ^3 and ℓ^4 bounds via interpolation in Lorentz scales).

2. Automated inequality search with Lorentz primitives. One practical benefit of replacing truncation lemmas by exact Lorentz norms is that a large class of arguments becomes a concatenation of a small set of reusable inequalities:

- dyadic layer-cake identities of the form $\|f\|_{p,1} \asymp \sum_k 2^k |\{f \sim 2^k\}|^{1/p}$,
- Lorentz-space Hölder templates (e.g. $\langle f, g\rangle \leq \|f\|_{p,1}\|g\|_{p',\infty}$),
- Young-type convolution bounds at endpoint scales,
- and bookkeeping rules such as subadditivity under disjoint unions at the C-level.

Once these are treated as primitives, the dependence on auxiliary cutoff parameters largely disappears; the remaining quantitative content is encoded in a graph of inequalities and exponent constraints. This makes the overall proof architecture amenable to computer assistance in two complementary ways.

First, one can attempt *verification*: given a human-written proof that stays within a prescribed library of Lorentz and convolution inequalities, a checker can confirm that each step is dimensionally consistent, that all exponents match, and that implicit constants are tracked correctly (at least qualitatively). This is not a replacement for mathematical insight, but it is a guard against the kinds of bookkeeping errors that become common when a proof contains many truncation thresholds and dyadic pigeonholes.

Second, one can attempt *optimisation*: many propagation arguments amount to choosing intermediate parameters (thresholds δ, sizes of symmetry sets, decomposition depths) to optimise the final exponent. In a Lorentz-based presentation, these parameters typically appear as exponents in inequalities rather than as ad hoc cutoffs, and thus the optimisation often reduces to a constrained minimisation problem (sometimes piecewise-linear after taking logs). It is plausible that one can recover, and perhaps improve, certain exponent choices by systematic search over admissible inequality chains, with the Lorentz quasi-norm serving as the correct endpoint that prevents the optimiser from exploiting illegal truncation artefacts.

3. Stability and inverse theorems for Lorentz control. The definition of $\mathsf{C}(A)$ is an operator norm, and operator norms invite an inverse theory: if $\mathsf{C}(A)$ is unusually large or unusually small, what must A look like, and what do near-extremising test functions f look like?

On the "large" side, one expects that $\mathsf{C}(A)$ close to its maximal value should force substantial additive structure. Indeed, testing $f = 1_A$ shows that $\mathsf{C}(A)$ controls $\|1_A * 1_A\|_3$, hence a third-moment statistic of additive representations within A. A natural conjectural statement is a stability version of Balog–Szemerédi–Gowers adapted to the L^3 certificate: if $\mathsf{C}(A) \geq \eta$ (equivalently $\kappa(A) \geq \eta^3$), then A contains a large subset $A' \subset A$ with small doubling, with polynomial dependence on η. Bloom's arguments already move in this direction; what is missing is a sharp understanding of near-extremisers, which would ideally upgrade polynomial losses to near-optimal ones and clarify the correct model examples.

On the "small" side, one may ask for structural interpretations of $\mathsf{C}(A) \ll |A|^{-\sigma}$ beyond the immediate consequences for energy and sumsets. For instance, for sets coming from geometry (convex sets, points on a curve, sets with few collinearities), $\mathsf{C}(A)$ is small because an incidence bound suppresses high-multiplicity additive coincidences. It would be valuable to formalise this as a general principle: small $\mathsf{C}(A)$ should be equivalent, in appropriate categories, to the non-existence of certain local concentration patterns. This would amount to an "inverse incidence theorem" phrased purely in additive terms.

Finally, because $\mathsf{C}(A)$ is defined by a supremum over weights $f \geq 0$, one can ask about *extremising weights*. Are there always near-extremisers supported on a small number of dyadic levels (or, conversely, can genuine extremisers require many levels, as suggested by the toy example)? If one could show that extremisers have additional rigidity (for instance, approximate indicator structure after normalisation), then the passage from weighted to unweighted statements would become more robust, and one might obtain cleaner extraction lemmas in the style of BSG without auxiliary regularisation.

In all three directions, the guiding idea is the same: once the endpoint space $\ell^{3/2,1}$ is accepted as the natural domain for tail-free control, the remaining difficulties are no longer about truncation bookkeeping but about genuine combinatorial or geometric input. This is precisely the setting in which one can hope for both conceptual clarity and quantitative improvement.