

VANISHING SUMS OF ROOTS OF UNITY AND THE FAVARD LENGTH OF SELF-SIMILAR PRODUCT SETS

IZABELLA LABA, CALEB MARSHALL

ABSTRACT. We improve a special case of the Lam-Leung lower bound [22] on the number of elements in a vanishing sum of N -th roots of unity. Using this result, we extend the Favard length estimates due to Bond, Laba, and Volberg [3] to a new class of rational product Cantor sets in \mathbb{R}^2 .

1. INTRODUCTION

A *vanishing sum of roots of unity* is an expression of the form

$$(1.1) \quad z_1 + \cdots + z_k = 0,$$

where z_1, \dots, z_k are N -th roots of unity for some $N \in \mathbb{N}$. In general, we do not require the z_j to be distinct or primitive N -th roots of unity. Vanishing sums of roots of unity have been studied extensively in number theory, see e.g., [14], [10], [12], [13], [22], [23], [28], [29], [30], [31]. Of particular interest is the following result, taken from [22].

Theorem 1.1. *Let $N = \prod_{j=1}^K p_j^{\alpha_j}$, where p_1, \dots, p_K are distinct primes and $\alpha_1, \dots, \alpha_K \in \mathbb{N}$. Suppose that (1.1) holds, where z_1, \dots, z_k are N -th roots of unity. Then there exist nonnegative integers a_1, \dots, a_K such that $k = \sum_{j=1}^K a_j p_j$. In particular, $k \geq \min\{p_j\}$.*

A key principle motivating our work is that (1.1) can be rephrased in terms of divisibility of polynomials by cyclotomic factors. Recall that the s -th **cyclotomic polynomial** $\Phi_s(X)$ is the unique monic, irreducible polynomial whose roots are the primitive s -th roots of unity. In other words, $z \in \mathbb{C}$ is a root of $\Phi_s(z) = 0$ if and only if we have $z = e^{2\pi i d/s}$ for some $d \in \mathbb{Z}$ with $(d, s) = 1$. Alternatively, Φ_s can be defined inductively via the identity

$$(1.2) \quad X^n - 1 = \prod_{s|n} \Phi_s(X).$$

Thus $\Phi_1(X) = X - 1$, and

$$(1.3) \quad 1 + X + X^2 + \cdots + X^{n-1} = \prod_{s|n, s \neq 1} \Phi_s(X) \text{ for } n > 1.$$

In particular, if p is a prime number, then $\Phi_p(X) = 1 + X + \cdots + X^{p-1}$. Thus

$$0 = \Phi_p(e^{2\pi i/p}) = 1 + e^{2\pi i/p} + \cdots + e^{2\pi i(p-1)/p}$$

provides an example of a vanishing sum of roots of unity. In general, however, such sums can be much more complicated.

Date: March 14, 2022.

Assuming that z_1, \dots, z_k are (not necessarily primitive) N -th roots of unity, we have $z_\ell = e^{2\pi i a_\ell / N}$ for some $a_\ell \in \mathbb{Z}$. Then (1.1) holds if and only if

$$(1.4) \quad A(e^{2\pi i / N}) = 0,$$

where $A(X) = \sum_{a=0}^{N-1} w(a)X^a$ and

$$w(a) = |\{\ell \in \{1, \dots, k\} : z_\ell = e^{2\pi i a_\ell / N} \text{ and } a_\ell \equiv a \pmod{N}\}|.$$

Since $\Phi_N(X)$ is the minimal polynomial of $e^{2\pi i / N}$, it follows that (1.4) holds if and only if $\Phi_N(X) | A(X)$.

The lower bound in Theorem 1.1 now takes the following form: suppose $\Phi_N(X) | A(X)$, where $A(X)$ is a polynomial with non-negative integer coefficients as above. Then

$$(1.5) \quad A(1) = \sum_{a=0}^{N-1} w(a) \geq \min_j p_j,$$

where p_1, \dots, p_K are the distinct prime factors of N .

We are interested in improvements to (1.5) when $A(X)$ is assumed to have multiple cyclotomic divisors. In general, such improvements are not possible without some additional assumptions on the cyclotomic factors of $A(X)$. To demonstrate, consider the simple example

$$A(X) = 1 + X^{q^n} + X^{2q^n} + \dots + X^{(p-1)q^n} = \Phi_p(X^{q^n}),$$

where $n \in \mathbb{N}$ and $p < q$ are distinct primes. For each $\alpha \in \{0, 1, \dots, n\}$, the number $e^{2\pi i q^{n-\alpha} / p}$ is a p -th primitive root of unity, therefore a root of $\Phi_p(X)$. Thus $A(e^{2\pi i / pq^\alpha}) = 0$, and consequently $\Phi_{pq^\alpha}(X) | A(X)$, for all $\alpha \in \{0, 1, \dots, n\}$; but we still only have $A(1) = p$. Since $n \in \mathbb{N}$ is arbitrary, we can find polynomials $A(X)$ such that $A(1)$ attains the lower bound of Theorem 1.1, but $A(X)$ has as many cyclotomic factors as we like.

However, we do have an improvement in the following case.

Theorem 1.2. *Let $A(X)$ be a polynomial with non-negative coefficients and distinct cyclotomic factors $\Phi_{m_1}(X), \dots, \Phi_{m_k}(X)$. Assume that there exist distinct prime numbers p, q , and exponents $\alpha_j, \beta_j \in \mathbb{N} \cup \{0\}$ such that $m_j = p^{\alpha_j} q^{\beta_j}$ for each $1 \leq j \leq k$. Assume further that $q \nmid A(1)$. Then we have the lower bound*

$$A(1) \geq p^{E_p}, \text{ where } E_p = |\{\alpha_1, \dots, \alpha_k\}|.$$

In words, E_p denotes the number of **distinct** exponents α_i appearing among the $m_j = p^{\alpha_j} q^{\beta_j}$.

A similar statement holds with p and q interchanged, so that if we assume that $p \nmid A(1)$ instead of $q \nmid A(1)$, we have an analogous lower bound in terms of q and its exponents. If neither of p, q divides $A(1)$, we may choose the maximum of the two lower bounds.

We note that the conclusion of Theorem 1.2 is immediate in the special case when $m_j = p^{\alpha_j}$ are all powers of the same prime p , with $\alpha_1, \dots, \alpha_k \geq 1$. Indeed, it is easy to deduce by induction from (1.3) that if p is prime and $\alpha \in \mathbb{N}$, then

$$(1.6) \quad \Phi_{p^\alpha}(X) = \Phi_p(X^{p^{\alpha-1}}) = 1 + X^{p^{\alpha-1}} + X^{2p^{\alpha-1}} + \dots + X^{(p-1)p^{\alpha-1}}, \quad \alpha \geq 1.$$

Hence $\Phi_{p^{\alpha_j}}(1) = p$, and $\prod_{j=1}^k \Phi_{p^{\alpha_j}}(X) | A(X)$ implies that $p^k = \prod_{j=1}^k \Phi_{p^{\alpha_j}}(1) | A(1)$. In particular, $p^k \leq A(1)$, as claimed. However, this argument does not apply to cyclotomic factors of the form $\Phi_{p^{\alpha_j} q^{\beta_j}}$ with $\alpha_j, \beta_j \geq 1$, since for such m_j we have $\Phi_{m_j}(1) = 1$.

Theorem 1.2 is motivated in part by its application to the Favard length problem for rational product Cantor sets. We give a brief introduction to this problem now, and state the relevant previous results, before presenting the extension derived from our Theorem 1.2.

Let $A, B \subset \mathbb{N}$ be finite sets with $\min(|A|, |B|) \geq 2$. With $L := |A||B|$, define the sets A_N and B_N recursively via the formulas $A_1 := L^{-1}A$, $A_{N+1} := A_N + L^{-N-1}A$, and similarly for B_N . Let

$$(1.7) \quad \mathcal{S}_N = A_N \times B_N + \{x \in \mathbb{R}^2 : |x| < L^{-N}\}, \quad N \in \mathbb{N}.$$

We may think of \mathcal{S}_N as the N -th iteration (i.e. approximation) of a self-similar set $\mathcal{S}_\infty \subset \mathbb{R}^2$, whose construction is as follows.

Let $\{z_j\}_{j=1}^L = A \times B \subset \mathbb{R}^2$; since $|A|, |B| \geq 2$, these points are distinct and non-colinear. Let $T_1, \dots, T_L : \mathbb{C} \rightarrow \mathbb{C}$ be the similarity maps $T_j(z) = \frac{1}{L}z + z_j$. We then define \mathcal{S}_∞ to be the unique compact set such that $\mathcal{S}_\infty = \bigcup_{j=1}^L T_j(\mathcal{S}_\infty)$. It is well known that such a set exists, has Hausdorff dimension at most 1 (equality follows if the **open set condition** holds, see [25]) and finite 1-dimensional Hausdorff measure.

For $\theta \in [0, \pi]$, let $proj_\theta : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the linear projection defined by $proj_\theta(x) = x_1 \cos \theta + x_2 \sin \theta$ for $x = (x_1, x_2) \in \mathbb{R}^2$. Since the z_j are not colinear, \mathcal{S}_∞ is unrectifiable, and it follows from a theorem of Besicovitch that $|proj_\theta(\mathcal{S}_\infty)| = 0$ for almost every $\theta \in [0, \pi]$ (see [25]). In particular, we must have the following ‘‘average’’ behaviour,

$$\lim_{N \rightarrow \infty} \frac{1}{\pi} \int_0^\pi |proj_\theta(\mathcal{S}'_N)| d\theta = 0,$$

where \mathcal{S}'_N is the L^{-N} -neighborhood of \mathcal{S}_∞ . We call the integrand in the above equality the **Favard length** of the set \mathcal{S}'_N . In particular, we must then have:

$$(1.8) \quad \lim_{N \rightarrow \infty} Fav(\mathcal{S}'_N) = Fav(\mathcal{S}_\infty) = 0,$$

While \mathcal{S}'_N is not necessarily equal to the finite scale set \mathcal{S}_N defined above, it can be covered by a bounded (independently of N) number of copies of \mathcal{S}_N , and vice versa. Therefore any quantitative estimates on $Fav(\mathcal{S}_N)$ and $Fav(\mathcal{S}'_N)$ are equivalent up to constants.

Although it is immediate that $Fav(\mathcal{S}_N) \rightarrow 0$ as $N \rightarrow \infty$, a more subtle question concerns estimating the rate of decay, both from above and from below. There has been significant interest in this issue, with Favard length estimates proved for various types of fractal sets in [1], [2], [3], [4], [6], [7], [9], [19], [26], [33], [35]. Motivation and an overview is provided in the review articles [17], [27], [34]; see also [5], [8], [11] for results of this type for curvilinear projections.

The following theorem, proved by Bond, Volberg, and the first author in [3], gives an asymptotic estimate for $Fav(\mathcal{S}_N)$ when A and B have small cardinality.

Theorem 1.3. *Define \mathcal{S}_N as in (1.7). If $|A|, |B| \leq 6$, then $Fav(\mathcal{S}_N) \lesssim N^{-\epsilon/\log \log N}$ for some $\epsilon > 0$.*

Here, the notation $Fav(\mathcal{S}_N) \lesssim N^{-\epsilon/\log \log N}$ means there exists a constant $C > 0$, independent of N , such that $Fav(\mathcal{S}_N) \leq CN^{-\epsilon/\log \log N}$ for all N large.

The proof in [3] relies upon the cyclotomic divisibility properties of certain polynomials associated to A and B . Specifically, define:

$$(1.9) \quad A(X) = \sum_{a \in A} X^a,$$

and similarly for B . We are assuming that $A, B \subset \mathbb{N}$, so that $A(X)$ and $B(X)$ are polynomials in $\mathbb{Z}[X]$. For our present purposes, it will suffice to consider A and B separately; and so, we present these requirements for the polynomial $A(X)$. Recalling that $L = |A||B|$, we define the following factors of $A(X)$.

Definition 1.4. *We have $A(X) = \prod_{i=1}^4 A^{(i)}(X)$, where each $A^{(i)}(X)$ is a product of the irreducible factors of $A(X)$ in $\mathbb{Z}[X]$, defined as follows (by convention, an empty product is identically equal to 1):*

- $A^{(1)}(X) = \prod_{s \in S_A^{(1)}} \Phi_s(X)$, $S_A^{(1)} = \{s \in \mathbb{N} : \Phi_s(X) | A(X), (s, L) \neq 1\}$,
- $A^{(2)}(X) = \prod_{s \in S_A^{(2)}} \Phi_s(X)$, $S_A^{(2)} = \{s \in \mathbb{N} : \Phi_s(X) | A(X), (s, L) = 1\}$,
- $A^{(3)}(X)$ is the product of those irreducible factors of $A(X)$ that have at least one root of the form $e^{2\pi i \xi_0}$, $\xi_0 \in \mathbb{R} \setminus \mathbb{Q}$,
- $A^{(4)}(X)$ has no roots on the unit circle.

We then define the **good** and **bad** factors A' and A'' of A :

$$(1.10) \quad A'(X) := A^{(1)}(X)A^{(3)}(X)A^{(4)}(X), \quad A''(X) := A^{(2)}(X).$$

Theorem 1.3 is then a consequence of the following more general result. We retain the notation of [3] below; however, we will later provide a notation better-tailored to our augmentation of the theorem.

Proposition 1.5. [3] *Let A, B be as in Theorem 1.3, but without the assumption that $\max(|A|, |B|) \leq 6$. Define*

$$(1.11) \quad S_A := \{s : \Phi_s | A \text{ and } (s, L) = 1\} = \{s : \Phi_s(X) | A''(X)\}, \quad s_A := \text{lcm}(S_A).$$

Suppose we can factor $s_A = s_{1,A}s_{2,A}$, with $s_{1,A}, s_{2,A} \in \mathbb{N}$ such that:

- $s_{2,A} < |A|$,
- $\Phi_s(X)$ does not divide $A(X)$ for any $s | s_{1,A}$.

Assume also that the same holds for B . Then $Fav(\mathcal{S}_N) \lesssim N^{-\epsilon / \log \log N}$ for some $\epsilon > 0$.

It is verified in [3] that the assumptions of Proposition 1.5 must hold if $\max(|A|, |B|) \leq 6$, hence Theorem 1.3 follows. The same proof also yields a stronger result in the case when $A^{(3)} = B^{(3)} = 1$.

Theorem 1.6. [3] *Let A, B be as in Theorem 1.3 or Proposition 1.5. Assume that all roots of $A(X)$ and $B(X)$ on the unit circle are roots of unity. Then $Fav(\mathcal{S}_N) \lesssim N^{-\epsilon}$ for some $\epsilon > 0$.*

Combining the framework of [3] with our Theorem 1.2 leads to the following result.

Theorem 1.7. *Define \mathcal{S}_N as in (1.7). Let $A(X) = \sum_{a \in A} X^a$, and similarly for $B(X)$. Define S_A and s_A as in (1.11), and similarly for B . Assume that s_A and s_B have at most two prime divisors each (not necessarily the same ones). Then*

$$\text{Fav}(\mathcal{S}_N) \lesssim N^{-\epsilon/\log \log N}$$

for some $\epsilon > 0$. If, in addition, all roots of $A(X)$ and $B(X)$ on the unit circle are roots of unity, we have the improvement $\text{Fav}(\mathcal{S}_N) \lesssim N^{-\epsilon}$ for some $\epsilon > 0$.

Remark 1.1. *The assumptions of Theorem 1.7 do not imply that the sets A, B satisfy the assumptions of Proposition 1.5; see Section 6.1 for an example.*

Remark 1.2. *It is not difficult to see that elements of S_A cannot be prime powers. Indeed, by (1.6) we have $\Phi_{p^\alpha}(1) = p$ if p is prime and $\alpha \in \mathbb{N}$. Hence $\Phi_{p^\alpha}(X) | A(X)$ implies that $p | A(1) = |A|$, and in particular that p^α cannot be relatively prime to $|A|$.*

It follows that if S_A is nonempty, then s_A must have at least two distinct prime factors. Thus Theorem 1.7 represents the simplest rational product case that goes beyond [3].

We are also able to increase the cardinality bound in Theorem 1.3.

Theorem 1.8. *Theorem 1.3 and Theorem 1.6 continue to hold with the assumption that $\max(|A|, |B|) \leq 6$ replaced by $\max(|A|, |B|) \leq 10$.*

The cardinality bound 10 in Theorem 1.8 is not a “hard” one, in the sense that the arguments used in the proof of the theorem continue to work for many larger sets. However, for sets of cardinality 11 and higher, the number of cases to consider increases significantly, making the proof more difficult to manage without contributing new ideas.

This article is organized as follows. The proof of Theorem 1.2 is given in Section 5, with the notation and basic cyclotomic divisibility tools provided in Section 4. This part of the paper is self-contained and does not involve any Favard length considerations, so that the reader interested only in Theorem 1.2 can work with just these two sections.

The proof of Theorem 1.7 consists of several steps. In Section 2, we explain how to combine the arguments of [3] (specifically, the *Set of Large Values* approach) with our Theorem 1.2 and our main intermediate result, Proposition 2.1, to prove the theorem. The proof of Proposition 2.1, given in Section 3, contains the main new ideas of this paper with regard to Favard length estimates. Specifically, while we follow the general approach of [3], we are able to use it more efficiently by splitting up the cyclotomic divisors of $A(X)$ and $B(X)$ into appropriate “clusters” and reorganizing the Set of Large Values construction accordingly.

An important part of the proof of Theorem 1.7 is that the lower bounds on $A(1)$ depending on the cyclotomic divisors of $A(X)$ (such as those in Theorem 1.2) have to match the size of a Set of Large Values that we can identify. In Section 6, we provide examples illustrating this. We also discuss briefly the barriers to extending our results to the case when s_A or s_B have 3 or more prime factors. Essentially, while certain simple examples of this type can be handled using the cuboid arguments in Section 6.4, a systematic approach to the general case would require additional new ideas. We conclude the paper with the proof of Theorem 1.8 in Section 7.

2. PROOF OF THEOREM 1.7

Our argument follows the general approach of [3], but with an additional argument to allow an application of Theorem 1.2. Most of the proof in [3] applies to general rational product sets; the only part that requires either the restriction $\max(|A|, |B|) \leq 6$ or the cyclotomic divisibility assumption in Proposition 1.5 is the SLV (**Set of Large Values**) argument in Sections 5 and 6. Below, we provide a short summary of what is needed, then explain how to modify this step in our two-prime setting. The proofs of the intermediate results are postponed until later sections.

We first state the SLV result we need. Define $\phi''_A(\xi) = A''(e^{2\pi i \xi})$, where $A''(X)$ is the “bad” factor associated to $A(X)$ in Definition 1.4; the function ϕ''_B is similarly defined. To extend the proof in [3] for sets A and B as in Theorem 1.7, it suffices to prove that the following holds under the assumptions of the theorem: given $t \in [\frac{1}{2}, 1]$, and given a large number m depending on N (in [3], m is at most logarithmic in N), there exists a Borel set $\Gamma = \Gamma(t, m) \subset [0, 1]$ satisfying the estimates

$$(2.1) \quad \Gamma - \Gamma \subset \left\{ \xi : \left| \prod_{k=0}^{m-1} \phi''_A(L^k \xi) \cdot \phi''_B(L^k t \xi) \right| \geq L^{-C_1 m} \right\},$$

$$(2.2) \quad |\Gamma| \geq C_2 L^{-(1-\epsilon)m},$$

where $C_1, C_2, \epsilon > 0$ are constants independent of N and m . The number $t = \tan(\theta)$ depends on the direction of an appropriately selected one-dimensional projection of S_N . We refer to any set $\Gamma \subset [0, 1]$ satisfying the above as an **SLV set** for the function $\phi''_t(\xi) = \phi''_A(\xi)\phi''_B(t\xi)$. Thus, to prove Theorem 1.7, it suffices to construct an SLV set $\Gamma \subset [0, 1]$ associated to sets $A, B \subset \mathbb{N}$ satisfying the number theoretic conditions of Theorem 1.7.

For $A \subset \mathbb{N}$, define S_A and s_A as in (1.11), with $S_A =: \{s_1, \dots, s_J\}$. Let also

$$\Sigma_A := \{x \in [0, 1] : \Phi_s(e^{2\pi i x}) = 0 \text{ for some } s \in S_A\}.$$

Observe that we cannot have $\Phi_1|A$. Indeed, $\Phi_1(X) = X - 1$, and if this were a factor of $A(X)$, we would have $A(1) = 0$, a contradiction. In particular, we have $\{0, 1\} \cap \Sigma_A = \emptyset$.

We then have the following single-scale SLV result for A (see Section 3.3 for a proof).

Proposition 2.1. *Let $A \subset \mathbb{N}$, and let $s_A = \prod_{i=1}^K p_i^{n_i}$, where p_1, \dots, p_K are distinct primes. For each $s_j \in S_A$, let $s_j = \prod_{i=1}^K p_i^{\alpha_{i,j}}$ be its prime factor decomposition. Assume that there exists an index $i \in \{1, \dots, K\}$ such that $\alpha_{i,j} \geq 1$ (so that $p_i|s_j$) for all $j \in \{1, \dots, J\}$. Let $E_i := |\{\alpha_{i,1}, \dots, \alpha_{i,J}\}|$ (that is, E_i counts the number of **distinct** exponents appearing on the prime factor p_i among the $s_j \in S_A$). Then for any $0 < \lambda_i < p_i^{-1}$ there exists a 1-periodic set $\Gamma_A \subset \mathbb{R}$ satisfying*

$$(2.3) \quad \text{dist}(\Gamma_A - \Gamma_A, \Sigma) > 0,$$

$$(2.4) \quad |[0, 1] \cap \Gamma_A| > \lambda_i^{E_i}.$$

To deduce Theorem 1.7 from Proposition 2.1, we will use it to find a multiscale SLV set Γ avoiding small values of both ϕ''_A and ϕ''_B as in (2.1). Then we must check that the size of this set is bounded from below as in (2.2). This is the step where we need Theorem 1.2. The full argument is below.

Proof of Theorem 1.7. Begin with two sets $A, B \subset \mathbb{N}$ satisfying the assumptions of Theorem 1.7, and consider the associated sets $\Gamma_A, \Gamma_B \subset \mathbb{R}$ provided by Proposition 2.1. From equation (2.3), there are constants $c_A, c_B > 0$ such that

$$(2.5) \quad \phi_A''(\xi) > c_A, \quad \forall \xi \in \Gamma_A - \Gamma_A$$

and similarly for ϕ_B'' on $\Gamma_B - \Gamma_B$. Setting $\Gamma_{k,A} = L^{-k} \cdot \Gamma_A$, we then have $\phi_A''(L^k \xi) > c_A$ on $\Gamma_{k,A}$, and similarly $\phi_B''(L^k \xi) > c_B$ on $\Gamma_{k,B} = L^{-k} \cdot \Gamma_B$. Let also

$$\nu_A := |[0, 1] \cap \Gamma_A|$$

be the density of Γ_A in $[0, 1]$, and similarly for ν_B .

Fix a large integer $M > 0$. The same pigeonholing argument as in [3, Proposition 5.1] (see also the proof of Lemma 3.4 in this article) furnishes translation parameters $\tau_{k,A}, \tau_{k,B} \in [0, M]$ so that the set

$$(2.6) \quad \Gamma := \bigcap_{k=0}^{m-1} (\Gamma_{k,A} - \tau_{k,A}) \cap \bigcap_{k=0}^{m-1} (t^{-1}\Gamma_{k,B} - \tau_{k,B})$$

satisfies the inequality

$$(2.7) \quad |\Gamma \cap [0, 1]| \geq \left(\frac{(M-1)(M-t^{-1})}{M^2} \nu_A \nu_B \right)^m.$$

Since the integers of S_A (and also S_B) share only two prime factors, the lower bound of Theorem 1.2 applies to $|A|$ and $|B|$. Combining this with equation (2.4) of Proposition 2.1 gives

$$\nu_A > \frac{1}{|A|}, \quad \nu_B > \frac{1}{|B|}.$$

Since these inequalities are strict, if $M > 0$ was chosen large-enough, there is an $\epsilon > 0$ with

$$\frac{(M-1)(M-t^{-1})}{M^2} \nu_A \nu_B \geq L^{-(1-\epsilon)}.$$

Note that our choice of $M = M(m, t)$ may depend upon m (therefore N) and t ; however, this does not affect the rest of the argument.

Thus, the set Γ defined in (2.6) satisfies (2.2). Moreover, since

$$\Gamma_{k,A} - \Gamma_{k,A} = (\Gamma_{k,A} - \tau_{k,A}) - (\Gamma_{k,A} - \tau_{k,A}),$$

and similarly for $\Gamma_{k,B} - \Gamma_{k,B}$, for each $0 \leq k \leq m-1$ we have

$$\phi_A''(L^k \xi), \phi_B''(tL^k \xi) > c, \quad \forall \xi \in \Gamma - \Gamma,$$

where $c = \min\{c_A, c_B\}$. Hence, (2.1) holds with $C_1 = \frac{2 \log(1/c)}{\log L}$, and our choice of $C_1 > 0$ is independent of N . Thus, the set $[0, 1] \cap \Gamma$, with Γ given by (2.6), is an SLV-set for the function $\phi_A''(x)\phi_B''(tx)$, as required. \square

3. CONSTRUCTING SINGLE-SCALE SETS Γ

The main new idea in the proof of Proposition 2.1 is the following “cluster reduction”. Let $A \subset \mathbb{N}$ be a finite set, and define S_A and s_A as in (1.11). Then we may split S_A into subsets called **clusters**, construct an SLV set associated to each cluster, and then take the intersections of appropriate translates of them to get the set Γ in the proposition.

The results of Section 3.1 and 3.2 apply to any finite set $A \subset \mathbb{N}$ and any splitting of A into clusters. In Section 3.1, we follow a number-theoretic approach due to Matthew Bond and the first author (cf. [17, Conjecture 4.6]), which extends slightly that of [3]. We should note here that Conjecture 4.6 in [17] turns out to be false, with a counterexample provided here [refer to example]. Therefore, if we simply tried to use the construction in Section 3.1 with $A = \mathcal{C}$ as a single cluster, our quantitative bounds on Γ would not be good enough. However, we can use the same construction more efficiently if we split up A into smaller clusters, construct a cluster-dependent set $\Gamma(\mathcal{C})$ for each one, then take the intersection of appropriate translates of the sets thus obtained.

In Section 3.3, we set up the cluster splitting that provides the requisite quantitative estimate (2.4) in Proposition 2.1. This part requires the additional assumption (stated in the proposition) on the prime factorization of the elements of S_A .

3.1. A single-cluster SLV set. Let s_1, \dots, s_J be an enumeration of the distinct elements of S_A .

Definition 3.1. A subset $\mathcal{C} \subset S_A$ is called a **cyclotomic divisor cluster**, or (for short) a **cluster**, of A .

Fix a cluster $\mathcal{C} \subset S_A$. Relabelling the elements of S_A if necessary, we may assume that $\mathcal{C} = \{s_1, \dots, s_I\}$ for some $I \leq J$. Let $N = \text{lcm}(s_1, \dots, s_I)$. Suppose that we can write $N = QU$, where

$$(3.1) \quad s_j \text{ does not divide } Q \text{ for any } j \in \{1, \dots, I\}.$$

In particular, it follows that $(s_j, U) > 1$ for each $s_j \in \mathcal{C}$. It will be to our advantage to choose Q as large as possible.

Each $s_j \in \mathcal{C}$ can be written as $s_j = r_j t_j$, where $r_j := (s_j, Q)$ and $t_j := s_j / (s_j, Q)$. Let $T := \max(t_1, \dots, t_I)$. Define

$$(3.2) \quad \Gamma(\mathcal{C}, \rho) = \left\{ \xi \in \mathbb{R} : \text{dist} \left(\xi, \frac{1}{Q} \mathbb{Z} \right) < \frac{\rho}{2} \right\},$$

where $0 < \rho < (QT)^{-1}$. The next two lemmas guarantee that the set in (3.2) has the properties we need. Specifically, the required arithmetic structure of $\Gamma(\mathcal{C}, \rho)$ is verified in Lemma 3.2, and Lemma 3.3 furnishes a lower bound on $|[0, 1] \cap \Gamma(\mathcal{C}, \rho)|$.

Lemma 3.2. *The set $\Gamma := \Gamma(\mathcal{C}, \rho)$ defined above satisfies*

$$\text{dist}(\Gamma - \Gamma, \Sigma(\mathcal{C})) > 0$$

where

$$(3.3) \quad \Sigma(\mathcal{C}) := \{ \xi \in \mathbb{R} : \Phi_{s_j}(e^{2\pi i \xi}) = 0 \text{ for some } s_j \in \mathcal{C} \}.$$

Proof. It suffices to prove that $\Gamma - \Gamma$ is disjoint from $\Sigma(\mathcal{C})$. The conclusion then follows by shrinking ρ a little bit more if needed.

Let $\xi \in \Sigma(\mathcal{C})$. Then there is an $s_j \in \mathcal{C}$ such that $\Phi_{s_j}(e^{2\pi i\xi}) = 0$, so that $\xi = b/s_j$ for some $b \in \mathbb{Z}$ with $(b, s_j) = 1$. Suppose $\xi \in \Gamma - \Gamma$, then there is an integer a such that

$$\left| \frac{a}{Q} - \frac{b}{s_j} \right| = \left| \frac{a}{Q} - \frac{b}{t_j r_j} \right| < \frac{1}{QT}.$$

Multiply this by Qt_j :

$$\left| at_j - \frac{bQ}{r_j} \right| < \frac{t_j}{T} < 1.$$

But $r_j|Q$, so that bQ/r_j is integer. Therefore

$$at_j = \frac{bQ}{r_j},$$

and in particular t_j divides bQ/r_j . Since $(b, s_j) = 1$, we also have $(b, t_j) = 1$, so that t_j divides Q/r_j . But then $s_j = t_j r_j$ divides Q , contradicting (3.1). \square

Lemma 3.3. *Let $\Gamma = \Gamma(\mathcal{C}, \rho)$ be as above, and let $\lambda < T^{-1}$. Then Γ is 1-periodic, and there exists a choice of $0 < \rho < (QT)^{-1}$ such that*

$$(3.4) \quad |[0, 1] \cap \Gamma| > \lambda.$$

Proof. The periodicity is clear from the definition. Let $\rho \nearrow (QT)^{-1}$, then

$$|[0, 1] \cap \Gamma| = Q\rho \nearrow \frac{Q}{QT} = \frac{1}{T}.$$

Thus it suffices to choose ρ sufficiently close to $(QT)^{-1}$. \square

3.2. Combining several clusters.

Lemma 3.4. *Suppose that $\mathcal{C}^1, \dots, \mathcal{C}^k$ are clusters associated to some $A \subset \mathbb{N}$. For each $l \in \{1, \dots, k\}$, let $\Gamma^l := \Gamma(\mathcal{C}^l, \rho^l)$ be the set defined in (3.2) and satisfying (3.4), with the corresponding parameters $Q_l, T_l, \rho_l, \lambda_l$. Then there exist translation parameters $\tau_1, \dots, \tau_k \in [0, 1]$ such that:*

$$(3.5) \quad \left| [0, 1] \cap \bigcap_{l=1}^k (\Gamma^l + \tau_l) \right| > \prod_{l=1}^k \lambda_l,$$

Furthermore, if we define $\Gamma^{1, \dots, k} := \bigcap_{l=1}^k (\Gamma^l + \tau_l)$ with this choice of τ_l , then

$$(3.6) \quad \text{dist}(\Gamma^{1, \dots, k} - \Gamma^{1, \dots, k}, \bigcup_{l=1}^k \Sigma(\mathcal{C}^l)) > 0$$

where $\Sigma(\mathcal{C}^l)$ is defined as in (3.3) with $\mathcal{C} = \mathcal{C}^l$.

Proof. We first note that for any fixed l ,

$$\text{dist}(\Gamma^{1, \dots, k} - \Gamma^{1, \dots, k}, \Sigma(\mathcal{C}^l)) \geq \text{dist}(\Gamma^l - \Gamma^l, \Sigma(\mathcal{C}^l)) > 0$$

by Lemma 3.2. Hence (3.6) holds for any choice of the parameters τ_l .

We now prove (3.5). The proof uses essentially the same argument as the proof of [3, Proposition 5.1], except that the large parameter M is not needed since all the sets Γ^l are 1-periodic. For each $x \in [0, 1)$, consider the function

$$\Psi(x) := \int_{[0,1]^k} \prod_{l=1}^k \mathbf{1}_{\Gamma^l}(x + \tau_l) d\tau_1 \cdots d\tau_k,$$

where $\tau_l \in [0, 1]$ are independent translation parameters. Using this independence, we evaluate $\Psi(x)$ as a product of single-variable averages. We have

$$\int_0^1 \mathbf{1}_{\Gamma^l}(x + \tau_l) d\tau_l = |[0, 1] \cap (\Gamma^l - x)| > \lambda_l,$$

by (3.4) and the 1-periodicity of each Γ^l . This leads to the pointwise lower bound

$$\Psi(x) > \prod_l \lambda_l.$$

Integrating and applying Fubini's theorem, we get that

$$\prod_l \lambda_l < \int_0^1 \Psi(x) dx = \int_{[0,1]^k} \left| [0, 1] \cap \bigcap_l (\Gamma^l - \tau_l) \right| d\tau_1 \cdots d\tau_k,$$

In particular, there exist translation parameters $\tau_1, \dots, \tau_k \in [0, 1]$ so that

$$\left| [0, 1] \cap \bigcap_l (\Gamma^l - \tau_l) \right| > \prod_l \lambda_l,$$

as claimed. \square

3.3. Choosing the clusters. We divide the set S_A into clusters based on their number-theoretic properties. Lemma 3.4 then produces an SLV set associated to these clusters, with an appropriate lower bound on its measure. We first introduce notation that allows us to partition the set $S_A := \{s_1, \dots, s_J\}$ in a useful way.

Each $s_j \in S_A$ has the form $s_j = p_1^{\alpha_{1,j}} \cdot p_2^{\alpha_{2,j}} \cdots p_K^{\alpha_{K,j}}$ where p_1, \dots, p_K are distinct prime numbers. For each $1 \leq i \leq K$, let

$$\text{EXP}(i) := \{\alpha_{i,1}, \dots, \alpha_{i,J}\}.$$

In words, $\text{EXP}(i)$ contains the exponents appearing on p_i among the $s_j \in S_A$. Observe that some of the numbers $\alpha_{i,1}, \dots, \alpha_{i,J}$ may be equal. We set $\#\text{EXP}(i) = E_i$, the number of **distinct** exponents in $\text{EXP}(i)$.

Definition 3.5. Fix $1 \leq i \leq K$. Then, for each $\alpha \in \text{EXP}(i)$, define

$$\mathcal{C}^{i,\alpha} = \{s_j \in S_A : s_j = p_i^\alpha \cdot q, \text{ for some } q \in \mathbb{N} \text{ with } (p_i, q) = 1\}$$

That is, $\mathcal{C}^{i,\alpha}$ is the cluster of s_j such that $\alpha_{i,j} = \alpha$.

We have the following lemma concerning clusters associated to non-zero $\alpha \in \text{EXP}(i)$.

Lemma 3.6. Fix $1 \leq i \leq K$. Let $\alpha \in \text{EXP}(i) \setminus \{0\}$, and let $\mathcal{C} := \mathcal{C}^{i,\alpha}$. Let $0 < \lambda_i < p_i^{-1}$. Then there is a choice of $Q = Q^{i,\alpha}$ and $\rho = \rho^{i,\alpha}$ such that the set $\Gamma_{i,\alpha} := \Gamma(\mathcal{C}^{i,\alpha}, \rho^{i,\alpha})$ constructed in Section 3.1 satisfies the estimate

$$|[0, 1] \cap \Gamma_{i,\alpha}| > \lambda_i.$$

Proof. Since i and α are fixed throughout the proof, we suppress them for now and use the notation of Section 3.1 with $\mathcal{C} = \mathcal{C}^{i,\alpha}$.

Relabelling the elements of S_A if necessary, we may assume that $\mathcal{C} = \{s_1, \dots, s_I\}$ for some $I \leq J$. Each $s_j \in \mathcal{C}$ can be written as $s_j = p_i^\alpha q_j$, where $(q_j, p_i) = 1$. We then define

$$Q := p_i^{\alpha-1} \operatorname{lcm}(q_1, \dots, q_I).$$

Then for each $j \in \{1, \dots, I\}$ we have $(s_j, Q) = p_i^{\alpha-1} q_j$ and

$$t_j = \frac{s_j}{(s_j, Q)} = p_i.$$

It follows that we can take $T = \max(t_j) = p_i$.

Let $\Gamma_{i,\alpha} := \Gamma(\mathcal{C}, \rho)$ be the set defined in (3.2). Then we may choose $0 < \rho < (QT)^{-1}$ so that the conclusions of Lemmas 3.2 and 3.3 hold with $\lambda = \lambda_i$, as claimed. \square

Proof of Proposition 2.1. We now assume that $A \subset \mathbb{N}$ satisfies the hypotheses of the proposition. In particular, we have $0 \notin \operatorname{EXP}(i)$; this is simply a rephrasing of the requirement that $\alpha_{i,j} \geq 1$.

Let $\lambda_i < p_i^{-1}$. For each $\alpha \in \operatorname{EXP}(i)$, let $\Gamma_{i,\alpha}$ be the set provided by Lemma 3.6. Applying Lemma 3.4, we find translation parameters z_α so that

$$(3.7) \quad |[0, 1] \cap \bigcap_{\alpha \in \operatorname{EXP}(i)} (\Gamma_{i,\alpha} + z_\alpha)| > \lambda_i^{E_i},$$

Let $\Gamma_A := \bigcap_{\alpha} (\Gamma_{i,\alpha} + z_\alpha)$; then (3.7) shows that Γ_A has the correct size. Moreover, since $\Sigma_A = \bigcup_{\alpha \in \operatorname{EXP}(i)} \Sigma(\mathcal{C}^{i,\alpha})$, by (3.6) we have

$$\operatorname{dist}(\Gamma_A - \Gamma_A, \Sigma_A) > 0,$$

where $\Sigma(\mathcal{C}^{i,\alpha})$ is defined as in (3.3) with $\mathcal{C} = \mathcal{C}^{i,\alpha}$. This proves the proposition. \square

Remark 3.1. *The cluster splitting above is sufficient for our purposes if $K = 2$. In this case, since each element of S_A must have at least two distinct prime factors (cf. Remark 1.2), we must have $\alpha_{i,j} \geq 1$ for each $i \in \{1, 2\}$ and each j . Thus the construction above works with both choices of $i \in \{1, 2\}$. If $K \geq 3$, we would not be able to assume that, but we could still construct an SLV set by splitting the elements of S_A into disjoint sets $S_{A,i}$ such that $p_i | s$ for each $s \in S_{A,i}$, applying the construction in Section 3.3 to each such subset, and then proceeding as in Lemma 3.4 to take the intersection of appropriate translates of the sets thus obtained. We expect that optimizing over constructions of this type should generate SLV sets that approach the maximal allowed size. However, in the general case, we do not know how to prove matching lower bounds on the size of A . See Section 6.4 for further discussion.*

4. CYCLOTOMIC DIVISIBILITY TOOLS

The remainder of this article focuses on proving Theorem 1.2, our result on cyclotomic factor decompositions and vanishing sums of roots of unity. Some of the notation here has been borrowed from [18] and adapted to our setting.

4.1. Multisets. We will work in the ambient group \mathbb{Z}_M , where $M = p_1^{n_1} \cdots p_K^{n_K}$, p_1, \dots, p_K are distinct primes, and $n_1, \dots, n_K \in \mathbb{N}$. We will use m and N (possibly with subscripts) to denote divisors of M . For the purpose of proving Theorem 1.2, it would be sufficient to consider the case $K = 2$. The discussion in Section 6 will require the more general notation.

We use $A(X)$, $B(X)$, etc. to denote polynomials modulo $X^M - 1$ with integer coefficients. Each such polynomial $A(X) = \sum_{a \in \mathbb{Z}_M} w_A(a)X^a$ is associated with a weighted multiset in \mathbb{Z}_M , which we will also denote by A , with weights $w_A(x)$ assigned to each $x \in \mathbb{Z}_M$. (If the coefficient of X^x in $A(X)$ is 0, we set $w_A(x) = 0$.) In particular, if A has $\{0, 1\}$ coefficients, then w_A is the characteristic function of a set $A \subset \mathbb{Z}_M$. We will use $\mathcal{M}(\mathbb{Z}_M)$ to denote the family of all weighted multisets in \mathbb{Z}_M , and reserve the notation $A \subset \mathbb{Z}_M$ for sets. We will also use \mathcal{M}^+ to denote the family of all weighted multisets in \mathbb{Z}_M , i.e.,

$$\mathcal{M}^+ = \{A \in \mathcal{M}(\mathbb{Z}_M) : w_A(a) \geq 0 \text{ for all } a \in \mathbb{Z}_M\}.$$

Let $A \in \mathcal{M}^+(\mathbb{Z}_M)$, with the corresponding mask polynomial $A(X)$. We use $|A|$ to denote the cardinality of A with multiplicity, so that $|A| = A(1) = \sum_{x \in \mathbb{Z}_M} w_A(x)$. If $\Lambda \subset \mathbb{Z}_M$, we use $A \cap \Lambda$ to denote intersection with multiplicity, so that $w_{A \cap \Lambda}(x) = w_A(x)w_\Lambda(x)$.

If $N|M$, then any $A \in \mathcal{M}(\mathbb{Z}_M)$ induces a weighted multiset $A \bmod N$ in \mathbb{Z}_N , with the corresponding mask polynomial $A(X) \bmod (X^N - 1)$ and induced weights

$$(4.1) \quad w_A^N(x) = \sum_{x' \in \mathbb{Z}_M : x' \equiv x \pmod{N}} w_A(x'), \quad x \in \mathbb{Z}_N.$$

We extend the multiset notation to \mathbb{Z}_N , so that for example $\mathcal{M}(\mathbb{Z}_N)$ and $\mathcal{M}^+(\mathbb{Z}_N)$ denote the appropriate families of multisets.

We use convolution notation $A * B$ to denote the weighted sumset of $A, B \in \mathcal{M}(\mathbb{Z}_M)$:

$$(A * B)(X) = A(X)B(X), \quad w_{A * B}(x) = (w_A * w_B)(x) = \sum_{y \in \mathbb{Z}_M} w_A(x - y)w_B(y).$$

If one of the sets is a singleton, say $A = \{x\}$, we write $x * B = \{x\} * B$.

4.2. Grids and fibers. For $D|N|M$, a D -grid in \mathbb{Z}_N is a set of the form

$$\Lambda^N(x, D) := x + D\mathbb{Z}_N = \{x' \in \mathbb{Z}_N : D|(x - x')\}$$

for some $x \in \mathbb{Z}_N$. If $N = \prod_{j=1}^K p_j^{\alpha_j}$ is a divisor of M , with $0 \leq \alpha_j \leq n_j$, we let

$$D(N) := \prod_{j=1}^K p_j^{\gamma_j}, \quad \text{where } \gamma_j = \max(0, \alpha_j - 1) \text{ for } j = 1, \dots, K.$$

An N -fiber in the p_i direction is a set of the form $x * F_i^N \subset \mathbb{Z}_N$, where $x \in \mathbb{Z}_N$ and

$$(4.2) \quad F_i^N = \{0, N/p_i, 2N/p_i, \dots, (p_i - 1)N/p_i\}.$$

Thus $x * F_i^N = \Lambda^N(x, N/p_i)$.

4.3. Cuboids and structure results. As before, we work in \mathbb{Z}_M , where $M = \prod_{i=1}^K p_i^{n_i}$, and let $A \in \mathcal{M}(\mathbb{Z}_M)$. We will use the following notation from [18]. For multisets $\Delta \in \mathcal{M}(\mathbb{Z}_N)$, where $N|M$, we define the Δ -evaluations of A in \mathbb{Z}_N :

$$(4.3) \quad \mathbb{A}^N[\Delta] = \sum_{x \in \mathbb{Z}_N} w_A^N(x) w_\Delta^N(x).$$

The following special case is of particular interest.

Definition 4.1. Let M and N be as above, so that $N = \prod_{i=1}^K p_i^{n_i - \alpha_i}$, with $0 \leq \alpha_i \leq n_i$ for each $i = 1, \dots, K$. An N -**cuboid** is a multiset $\Delta \in \mathcal{M}(\mathbb{Z}_N)$ associated to a mask polynomial of the form

$$(4.4) \quad \Delta(X) = X^c \prod_{j: p_j | N} (1 - X^{d_j N / p_j})$$

with $(d_j, p_j) = 1$ for all j .

The geometric interpretation of N -cuboids Δ is as follows. With notation as in Definition 4.1, recall that $D(N) = N / \prod_{j: p_j | N} p_j$. Then the ‘‘vertices’’ $x \in \mathbb{Z}_N$ with $w_\Delta^N(x) \neq 0$ form a full-dimensional rectangular box in the grid $\Lambda(c, D(N))$, with one vertex at c and alternating ± 1 weights.

The following cyclotomic divisibility test has been known and used previously in the literature. The equivalence between (i) and (iii) is the Bruijn-Rédei-Schoenberg theorem on the structure of vanishing sums of roots of unity (see [14], [22], [23], [29], [30], [31]). For the equivalence (i) \Leftrightarrow (ii), see e.g. [32, Section 3], [16, Section 3].

Proposition 4.2. Let $A \in \mathcal{M}(\mathbb{Z}_N)$. Then the following are equivalent:

(i) $\Phi_N(X) | A(X)$,

(ii) For all N -cuboids Δ , we have

$$(4.5) \quad \mathbb{A}^N[\Delta] = 0,$$

(iii) $A \bmod N$ is a linear combination of N -fibers, so that

$$A(X) = \sum_{i: p_i | N} P_i(X) F_i^N(X) \pmod{X^N - 1},$$

where $P_i(X)$ have integer (but not necessarily nonnegative) coefficients.

Proposition 4.2 can be strengthened as follows if N has only two distinct prime factors. This goes back to the work of de Bruijn [14]; a self-contained proof is provided in [22, Theorem 3.3].

Lemma 4.3. Let $A \in \mathcal{M}^+(\mathbb{Z}_N)$. Assume that $\Phi_N | A$, where N has two distinct prime factors p_1, p_2 . Then $A \bmod N$ is a linear combination of N -fibers with nonnegative weights. In other words,

$$A(X) = P_1(X) F_1^N(X) + P_2(X) F_2^N(X) \pmod{X^N - 1},$$

where P_1, P_2 are polynomials with nonnegative coefficients.

The equivalence (i) \Leftrightarrow (ii) in Proposition 4.2 has the following consequence.

Lemma 4.4. *Assume that $N|M$ and $A \in \mathcal{M}(\mathbb{Z}_N)$. Then $\Phi_N|A$ if and only if $\Phi_N|(A \cap \Lambda)$ for every $D(N)$ -grid Λ .*

5. PROOF OF THEOREM 1.2

We are now ready to prove Theorem 1.2. For the reader's convenience, we state it here again in the notation of Section 4.

Proposition 5.1. *Assume that $K = 2$, and write $p = p_1, q = p_2$ for short. Let $A \in \mathcal{M}^+(\mathbb{Z}_M)$, and let m_1, m_2, \dots, m_r be divisors of M such that $m_j = p^{\alpha_j} q^{\beta_j}$, where*

$$1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_r.$$

Assume that $\Phi_{m_1} \Phi_{m_2} \dots \Phi_{m_r}|A$, and that $q \nmid |A|$. Then $|A| \geq p^r$.

Proof. We proceed by induction in r . For the base case, suppose that $r = 1$. By Lemma 4.3, $A \bmod m_1$ is a union of m_1 -fibers in the p and q directions. Since $q \nmid |A|$, at least one of these fibers must be in the p direction. Hence $|A| \geq p$.

Suppose now that $r \geq 2$, and that the proposition is true with r replaced by $r - 1$. We write \mathbb{Z}_M as a disjoint union of grids $\Lambda(y_i, p^{\alpha_1-1})$, where $y_0 := 0, y_1, \dots, y_{p^{\alpha_1-1}-1} \in \mathbb{Z}_M$. Let $A_i = A \cap \Lambda(y_i, p^{\alpha_1-1})$. By Lemma 4.4, we have $\Phi_{m_1} \Phi_{m_2} \dots \Phi_{m_r}|A_i$ for each i . Moreover, since $q \nmid |A|$, there exists at least one j such that $q \nmid |A_j|$. Without loss of generality, we may assume that $q \nmid |A_0|$. We will prove that $|A_0| \geq p^r$.

Write $m = m_1$ for short. Applying Lemma 4.3 to A_0 on the scale m , we see that $A_0 \bmod m$ is a union of m -fibers in the p and q directions. More explicitly,

$$A_0(X) = P_1(X)F_1^m(X) + P_2(X)F_2^m(X) \pmod{X^m - 1},$$

where P_1, P_2 are polynomials with nonnegative coefficients, F_1^m is an m -fiber in the $p = p_1$ direction, and F_2^m is an m -fiber in the $q = p_2$ direction. Moreover, since $A_0 \subset p^{\alpha_1-1}\mathbb{Z}_M$, we may choose $P_1(X)$ so that

$$P_1(X) = P_0(X^{p^{\alpha_1}}),$$

where $P_0(X)$ is a polynomial with nonnegative coefficients. Since $q \nmid |A_0|$ and $|F_2^m| = q$, it follows that $P_1(1)$ is not divisible by q .

We now split up A_0 further, as follows. Let $x_0 := 0, x_1, \dots, x_{p-1}$ be points in $\Lambda(0, p^{\alpha_1-1})$ such that $(x_i - x_j, M) = M/p^{\alpha_1-1}$ for $i \neq j$. We write $\Lambda(a, p^{\alpha_1-1}) = \bigcup_{j=0}^{p-1} \Lambda_j$, where $\Lambda_j = \Lambda(x_j, p^{\alpha_1})$, and let $A_{0,j} = A_0 \cap \Lambda_j$. Then for each $j = 0, 1, \dots, p-1$, we have

$$A_{0,j}(X) = P_1(X)X^{x_j} + Q_j(X)F_2^m(X) \pmod{X^m - 1},$$

where we used that any translated copy of F_2^m is either contained fully in Λ_j or disjoint from it. In particular, $|A_{0,j}|$ is not divisible by q . By another application of Lemma 4.4, we have $\Phi_{m_2} \dots \Phi_{m_r}|A_{0,j}$ for each j .

Applying the inductive assumption to $A_{0,j}$ for each j , we see that $|A_{0,j}| \geq p^{r-1}$. Hence

$$|A_0| = \sum_{j=0}^{p-1} |A_{0,j}| \geq p^r,$$

as claimed. □

6. EXAMPLES AND DISCUSSION

We discuss briefly a few motivating examples and the possibility of extending the results here to more general product sets. A minor inconvenience is that, in the Favard length setting, the set S_A defined in (1.11) depends on L , hence on both of the sets A and B in (1.7), and not just on A . Therefore, for the purpose of this discussion, consider the set

$$S'_A := \{s : \Phi_s | A \text{ and } (s, |A|) = 1\}.$$

Thus $S_A \subseteq S'_A$, with equality if $|A|$ and $|B|$ have the same prime factors. In the Favard length examples throughout the rest of this section, we will assume that $A = B$, so that $S_A = S'_A$.

6.1. Cluster splitting is necessary. It is not always advantageous to group divisors in large clusters. Consider the following example. Let $N = p^{10}q^{10}$, where p, q are distinct primes. Let

$$\begin{aligned} A(X) &= \frac{1 - X^N}{1 - X^{N/p}} \frac{1 - X^{N/p^9}}{1 - X^{N/p^{10}}} + \frac{1 - X^{N/q}}{1 - X^{N/q^9}} \frac{1 - X^{N/q^{10}}}{1 - X^{N/q^{10}}} \\ &= \left(1 + X^{N/p} + \dots + X^{(p-1)N/p}\right) \left(1 + X^{N/p^{10}} + \dots + X^{(p-1)N/p^{10}}\right) \\ &\quad + \left(1 + X^{N/q} + \dots + X^{(q-1)N/q}\right) \left(1 + X^{N/q^{10}} + \dots + X^{(q-1)N/q^{10}}\right). \end{aligned}$$

Then $\Phi_{s_1} \Phi_{s_2} | A(X)$ for $s_1 = N = p^{10}q^{10}$ and $s_2 = pq$, whereas $|A| = p^2 + q^2$.

Suppose we try to apply the construction of Section 3.1 (Lemmas 3.2 and 3.3) to the cluster $\mathcal{C} := \{s_1, s_2\}$. To do this, we need to choose an appropriate $Q|N$ satisfying (3.1). Since $s_2 = pq$ cannot divide Q , we must choose Q to be a power of only one of the primes. Let us be as generous as we can in the circumstances, and choose $Q = q^{10}$, where q is the larger prime. But then $(s_1, Q) = q^{10}$, so that $t_1 = p^{10}$. In order for that to be less than $|A|$, we would need $p^{10} < p^2 + q^2$, which is false if p and q are of about the same size. For the same reason, the set A does not satisfy the assumptions of Proposition 1.5.

However, the construction in Section 3.2 is more efficient. Consider the following cluster division. Let $\mathcal{C}^1 := \{p^{10}q^{10}\}$ and $\mathcal{C}^2 = \{pq\}$. We then choose $Q_1 = p^9q^{10}$ and $Q_2 = q$, with $T_1 = T_2 = p^{-1}$. Let Γ^1 and Γ^2 be the corresponding single-cluster SLV sets constructed in Section 3.1. By Lemma 3.4, there exist translation parameters $z_1, z_2 \in [0, 1]$ so that:

$$\left| [0, 1] \cap (\Gamma^1 + z_1) \cap (\Gamma^2 + z_2) \right| \geq \frac{1}{p^2} > \frac{1}{p^2 + q^2} = \frac{1}{|A|}.$$

This shows that dividing S_A into optimal clusters is a key component of obtaining sets Γ_A whose size compares favourably with $|A|$.

6.2. Explicit examples with two prime factors. The set A in Section 6.1 provides an explicit example of a set satisfying the two-prime assumption of Theorems 1.2 and 1.7. Other examples can be constructed in a similar way. For instance, let $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_k$, and let $N = p^\alpha q^\beta$, where p, q are distinct primes and $\alpha \geq \alpha_k$. Then the ‘‘long fiber’’ \mathcal{F} with the mask polynomial

$$\mathcal{F}(X) = \prod_{j=1}^k \Phi_p(X^{p^{\alpha_k - 1} q^\beta})$$

is divisible by all $\Phi_s(X)$ such that the exponent of p in the prime factorization of s is α_j for some $j \in \{1, \dots, k\}$. Translates of such fibers in both directions can be added to construct more complicated examples.

6.3. One scale, many primes. In the example in Section 6.1, we used two well separated scales ($p^{10}q^{10}$ and pq). If we allow s_A to have 4 or more distinct prime factors, then A may violate the assumptions of Proposition 1.5 in other ways.

The following example is due to Matthew Bond and the first author (unpublished). Let p_1, p_2, q_1, q_2 be distinct primes. Assume that $p_1 + q_1 = p_2 + q_2$, and, letting $N = p_1 + q_1$, that N is smaller than the product of any two distinct primes chosen from $\{p_1, p_2, q_1, q_2\}$. (For example, we could choose $p_1 = 3, p_2 = 5, q_1 = 13, q_2 = 11$, with $N = 16$.) Let also $s_1 = p_1q_1$ and $s_2 = p_2q_2$. Let A be a set of N integers such that

- $A \bmod s_1$ is a union of two s_1 -fibers, one in each direction, of cardinalities p_1 and q_1 ,
- $A \bmod s_2$ is a union of two s_2 -fibers, one in each direction, of cardinalities p_2 and q_2 .

This is easily produced via the Chinese remainder theorem. Then $\Phi_{s_1}(X)$ and $\Phi_{s_2}(X)$ divide $A(x)$. It follows that $p_1p_2q_1q_2$ divides s_A . We further note that none of $\Phi_{p_1}, \Phi_{p_2}, \Phi_{q_1}, \Phi_{q_2}$ divide $A(X)$.

Suppose that we have a factorization $s_A = s_{1,A}s_{2,A}$ satisfying the conditions of Proposition 1.5. If $s_{2,A}$ is divisible by at least two primes from $\{p_1, p_2, q_1, q_2\}$, then $s_{2,A} > N = |A|$, violating the first condition of the proposition. Therefore at most one of our four primes may divide $s_{2,A}$. Then, however, at least three of them must divide $s_{1,A}$. It follows that $s_{1,A}$ is divisible by at least one of s_1 and s_2 , violating the second condition.

Nonetheless, it turns out that the construction in Section 3.1 with the single cluster $\mathcal{C} = S_A = \{s_1, s_2\}$ is sufficient in this case. Indeed, let $Q = q_1q_2$, so that $t_1 = p_1$ and $t_2 = p_2$. Then $T = \max(p_1, p_2) < |A|$, and Lemmas 3.2 and 3.3 provide the requisite SLV set.

6.4. Lower bounds with more prime factors. Our current methods are not sufficient to extend Theorem 1.2 to the case when more than two distinct prime factors are allowed. Below, we indicate a single-step cuboid argument leading to a lower bound on $|A|$ in certain situations. This is enough to resolve simple examples such as those below. However, there does not seem to be any easy way to iterate the argument to allow more complicated configurations of cyclotomic divisors. Our proof of Theorem 1.2 fails at multiple points in this setting.

Lemma 6.1. *Let $A \in \mathcal{M}(\mathbb{Z}_M)$, where $M = \prod_{i=1}^K p_i^{n_i}$ and p_1, \dots, p_K are distinct primes. Suppose that $m|(M/p_i)$ for some $i \in \{1, \dots, K\}$. Assume further that $\Phi_{mp_i}|A$. Then $p_i|\mathbb{A}^m[\Delta]$ for any m -cuboid Δ .*

Proof. We consider A as a multiset in \mathbb{Z}_N , where $N = mp_i$. We define a family of weighted multisets $\Delta \in \mathcal{M}(\mathbb{Z}_N)$, as follows. Let $\mathfrak{J} = \{j \in \{1, \dots, K\} : p_j|m\}$. If $p_i \nmid m$, we consider Δ of the form

$$(6.1) \quad \Delta(X) = X^c \prod_{j \in \mathfrak{J}} (X^{d_j} - 1),$$

where $c \in \mathbb{Z}_N$ and $(d_j, N) = N/p_j$. If $p_i | m$, we instead consider

$$(6.2) \quad \Delta(X) = X^c (X^{d_i} - 1) \prod_{j \in \mathfrak{J}, j \neq i} (X^{d_j} - 1),$$

with c, d_j as above for $j \neq i$, and with $(d_i, N) = N/p_i^2$. In both cases, the induced multiset Δ in \mathbb{Z}_m is an m -cuboid, and any m -cuboid can be (non-uniquely) represented in this manner.

For $\nu = 0, 1, \dots, p_i - 1$, define $\Delta_\nu(X) = X^{\nu N/p_i} \Delta(x)$. We claim that

$$(6.3) \quad \mathbb{A}^m[\Delta] = \sum_{\nu=0}^{p_i-1} \mathbb{A}^N[\Delta_\nu].$$

Indeed, we have

$$\begin{aligned} \sum_{\nu=0}^{p_i-1} \mathbb{A}^N[\Delta_\nu] &= \sum_{\nu=0}^{p_i-1} \mathbb{A}^N[\nu N/p_i * \Delta_\nu] \\ &= \sum_{\nu=0}^{p_i-1} \sum_{x \in \mathbb{Z}_N} w_A^N(x - \nu N/p_i) w_\Delta^N(x) \\ &= \sum_{x \in \mathbb{Z}_m} w_A^m(x) w_\Delta^m(x) = \mathbb{A}^m[\Delta], \end{aligned}$$

as claimed.

Assume first that $p_i \nmid m$. Then $\Delta_\nu(X) - \Delta_{\nu'}(X)$ are N -cuboids for $\nu \neq \nu'$. Since $\Phi_N | A$, it follows from Proposition 4.2 that the corresponding cuboid evaluations are 0, hence $\mathbb{A}^N[\Delta_\nu]$, $\nu = 0, 1, \dots, p_i - 1$, are all equal. Thus $\mathbb{A}^m[\Delta] = p_i \mathbb{A}^N[\Delta_\nu]$ for any ν .

If $p_i | m$, the argument is only slightly more complicated. For each ν , we write $\Delta_\nu = \Delta_\nu^+ - \Delta_\nu^-$, where

$$\Delta_\nu^+(X) = X^{c+d_i+\nu N/p_i} \prod_{j \in \mathfrak{J}, j \neq i} (X^{d_j} - 1), \quad \Delta_\nu^-(X) = X^{c+\nu N/p_i} \prod_{j \in \mathfrak{J}, j \neq i} (X^{d_j} - 1).$$

Then $\Delta_\nu^+(X) - \Delta_{\nu'}^+(X)$ and $\Delta_\nu^-(X) - \Delta_{\nu'}^-(X)$ are N -cuboids for $\nu \neq \nu'$. Applying Proposition 4.2 as above, we see that $\mathbb{A}^N[\Delta_\nu^+]$, $\nu = 0, 1, \dots, p_i - 1$, are all equal, and similarly for $\mathbb{A}^N[\Delta_\nu^-]$. Using this together with (6.3), we get

$$\mathbb{A}^m[\Delta] = p_i (\mathbb{A}^N[\Delta_\nu^+] - \mathbb{A}^N[\Delta_\nu^-])$$

for any ν , proving the lemma in this case. \square

Lemma 6.2. *Let $M = \prod_{i=1}^K p_i^{n_i}$, where p_1, \dots, p_K are distinct primes. Let $\Lambda := \Lambda(c, p_i^\beta)$ for some $c \in \mathbb{Z}_M$, $1 \leq i \leq K$, and $0 < \beta < n_i$. Suppose that $A \in \mathcal{M}^+(\mathbb{Z}_M)$ is supported in Λ , in the sense that $w_A(x) = 0$ for all $x \notin \Lambda$. Let $M' = M/p_i^\beta$, and observe that the mapping*

$$\mathbb{Z}_{M'} \ni x \rightarrow c + p_i^\beta x \in \Lambda$$

is one-to-one. Define the ‘‘rescaled’’ multiset $A' \in \mathcal{M}(\mathbb{Z}_{M'})$ by

$$(6.4) \quad w_{A'}^{M'}(x) = w_A(c + p_i^\beta x).$$

Then for any m such that $p_i^\beta | m | M$, we have

$$\Phi_m | A \Leftrightarrow \Phi_{m/p^\beta} | A'.$$

Proof. Since the scaling in (6.4) maps m -cuboids supported in Λ to m' -cuboids in $\mathbb{Z}_{M'}$, the lemma follows from Proposition (4.2). \square

Proposition 6.3. *Let $A \in \mathcal{M}(\mathbb{Z}_M)$, where $M = \prod_{i=1}^K p_i^{n_i}$ and p_1, \dots, p_K are distinct primes. Suppose that there exist $i \in \{1, \dots, K\}$, $m_0 | (M/p_i^{n_i})$, and $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_\ell \leq n_i$ such that*

$$(6.5) \quad \Phi_{m_1} \dots \Phi_{m_\ell} | A, \quad \text{where } m_j = m_0 p_i^{\alpha_j}, \quad j = 1, \dots, \ell.$$

Then $p_i^\ell | \mathbb{A}^{m_0}[\Delta]$ for any m_0 -cuboid Δ .

Proof. The proof is by induction in ℓ . In order to streamline the proof, we note that the statement of the proposition holds trivially for $\ell = 0$, with no cyclotomic divisors assumed and the trivial conclusion $1 = p_i^0 | \mathbb{A}^{m_0}[\Delta]$. We will use this as the base case.

Assume now that $\ell \geq 1$, and that the proposition is true in any cyclic group $\mathbb{Z}_{M'}$ with ℓ replaced by $\ell - 1$. The inductive step is similar to the proof of Lemma 6.1. Assume that (6.5) holds. Let $m = m_1/p_i$ and $N = m_1 = m p_i$. Define

$$\mathfrak{J} := \{j \in \{1, \dots, K\} : p_j | m_0\}.$$

For any m_0 -cuboid Δ_0 , we may write $\mathbb{A}^{m_0}[\Delta_0]$ as a linear combination of expressions of the form $\mathbb{A}^m[\Delta]$, where

$$(6.6) \quad \Delta(X) = X^c \prod_{j \in \mathfrak{J}} (X^{d_j} - 1),$$

with $c \in \mathbb{Z}_N$ and $(d_j, N) = N/p_j$ for each j . We will prove that $p_i^\ell | \mathbb{A}^m[\Delta]$ for each such Δ .

Define

$$\Delta_\nu(X) = X^{c+\nu N/p_i} \prod_{j \in \mathfrak{J}, j \neq i} (X^{d_j} - 1)$$

for $\nu = 0, 1, \dots, p_i - 1$. Then

$$\mathbb{A}^m[\Delta] = \sum_{\nu=0}^{p_i-1} \mathbb{A}^N[\Delta_\nu].$$

Since $\Delta_\nu(X) - \Delta_{\nu'}(X)$ are N -cuboids for $\nu \neq \nu'$, Proposition 4.2 implies that $\mathbb{A}^N[\Delta_\nu]$ with $\nu = 0, 1, \dots, p_i - 1$ are all equal. Thus

$$(6.7) \quad \mathbb{A}^m[\Delta] = p_i \mathbb{A}^N[\Delta_\nu] \text{ for any } \nu.$$

For each $\nu = 0, 1, \dots, p_i - 1$, let $A_\nu := A \cap \Lambda(c + \nu N/p_i, p_i^{\alpha_1})$. Let also $M' := M/p_i^{\alpha_1}$, and define the rescaled multisets $A'_\nu \in \mathcal{M}(\mathbb{Z}_{M'})$ as in Lemma 6.2:

$$w_{A'_\nu}^{M'}(x) = w_A(c + \nu N/p_i + p_i^{\alpha_1} x).$$

By Lemmas 4.4 and 6.2, we have $\Phi_{m'_2} \dots \Phi_{m'_\ell} | A'_\nu$ for each ν , where $m'_j = m_0 p_i^{\alpha_j - \alpha_1}$. Furthermore, let Δ'_ν be the rescaling of Δ_ν , then Δ'_ν is an m_0 -cuboid in $\mathbb{Z}_{M'}$, and (with the obvious notation)

$$\mathbb{A}^N[\Delta_\nu] = (\mathbb{A}'_\nu)^{m_0}[\Delta'_\nu].$$

By the inductive assumption, the last quantity is divisible by $p_i^{\ell-1}$. The conclusion follows by combining this with (6.7). \square

Corollary 6.4. *Let $A \in \mathcal{M}^+(\mathbb{Z}_M)$, where $M = \prod_{i=1}^K p_i^{n_i}$ and p_1, \dots, p_K are distinct primes. Suppose that the assumptions of Proposition 6.3 are satisfied, and, additionally, there exists α_0 with $0 \leq \alpha_0 < \alpha_1$ such that*

$$\Phi_{m_*} \nmid A, \quad \text{where } m_* = m_0 p_i^{\alpha_0}.$$

Then $|A| \geq p_i^\ell$.

Proof. If $\alpha_0 = 0$, then there exists an m_0 -cuboid Δ such that $\mathbb{A}^{m_0}[\Delta] \neq 0$. Combining this with Proposition 6.3, we see that

$$|A| \geq |\mathbb{A}^{m_0}[\Delta]| \geq p_i^\ell,$$

as claimed.

If $\alpha_0 > 0$, we find instead an m_* -cuboid Δ such that $\mathbb{A}^{m_*}[\Delta] \neq 0$. As in the proof of Lemma 6.1, write $\Delta(X) = \Delta^+(X) - \Delta^-(X)$, where Δ^+, Δ^- are the faces of Δ perpendicular to the p_i direction, so that each of Δ^+, Δ^- must be contained in a single $p_i^{\alpha_0}$ -grid. At least one of $\mathbb{A}^{m_*}[\Delta^+]$ and $\mathbb{A}^{m_*}[\Delta^-]$ must be nonzero. Assume without loss of generality that $\mathbb{A}^{m_*}[\Delta^+] \neq 0$, and let Λ be the $p_i^{\alpha_0}$ -grid containing Δ^+ . Consider the restriction $A \cap \Lambda$ of A to Λ , and rescale it by a factor of $p_i^{-\alpha_0}$ as in Lemma 6.2. This reduces the proof to the case $\alpha_0 = 0$ as above. The details are left to the reader. \square

Example 6.1. Define \mathcal{S}_N as in (1.7), with $A = B$. Let $M := \text{lcm}(S_A) = \prod_{i=1}^K p_i^{n_i}$, where p_1, \dots, p_K are distinct primes. Let $m_0 | (M/p_i^{n_i})$. Suppose that

$$\{m_0 p_i^{\alpha_1}, m_0 p_i^{\alpha_2}, \dots, m_0 p_i^{\alpha_\ell}\} \subset S_A$$

for some $i \in \{1, \dots, K\}$ and $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_\ell \leq n_i$. Assume furthermore that, in the notation of Section 3.3,

$$\text{EXP}(i) := \{\alpha_1, \dots, \alpha_\ell\}.$$

(In words, no other powers of p_i appear in the prime factorization of elements of S_A .) By Corollary 6.4, we have $|A| \geq p_i^\ell$. Since $|A|$ is relatively prime to all elements of S_A , the inequality must be strict. Thus the SLV construction in Section 3.3, with the same choice of i , is sufficient in this case. It follows that \mathcal{S}_N satisfies the conclusions of Theorem 1.7 in this case.

The corollary below extends Lemma 6.1 in a different direction.

Corollary 6.5. *Let $A \in \mathcal{M}^+(\mathbb{Z}_M)$, where $M = \prod_{i=1}^K p_i^{n_i}$ and p_1, \dots, p_K are distinct primes. Suppose that $m p_1 \dots p_I | M$ for some $I < K$. Assume further that $\Phi_{p_i m} | A$ for $i = 1, \dots, I$, but $\Phi_m \nmid A$. Then $A(1) \geq p_1 \dots p_I$.*

Proof. By Lemma 6.1, we have $p_1 \dots p_I | \mathbb{A}^m[\Delta]$ for every m -cuboid Δ . Since $\Phi_m \nmid A$, by Proposition 4.2 there exists an m -cuboid Δ such that $\mathbb{A}^m[\Delta] \neq 0$. Hence for that Δ , we have

$$A(1) \geq |\mathbb{A}^m[\Delta]| \geq p_1 \dots p_I.$$

\square

Remark 6.1. *A stronger version of Corollary 6.5 can be obtained, with the same proof, by using Proposition 6.3 instead of Lemma 6.1. Since the statement would be significantly more complicated without contributing new ideas, we omit it here.*

Example 6.2. The following extends the example in Section 6.3. Let $A \in \mathcal{M}^+(\mathbb{Z}_M)$, with $M = p_1 \dots p_k Q$, where p_1, \dots, p_k are distinct primes and $(p_1 \dots p_k, Q) = 1$. Suppose that $S_A = \{s_1, \dots, s_k\}$, where for each s_j we have $p_j | s_j | p_j Q$.

We claim that $|A| > \max(p_1, \dots, p_k)$. Indeed, let $j \in \{1, \dots, k\}$. Since $\Phi_{s_j} | A$ and $\Phi_{s_j/p_j} \nmid A$, Corollary 6.5 with $m = s_j/p_j$ implies that $|A| \geq p_j$. Since we have $(|A|, s_j) = 1$ by the definition of S_A , the inequality must be strict.

For the purpose of an application to the Favard length problem, the single-cluster construction in Section 3.1 works for such A . Indeed, let $\mathcal{C} := S_A$, and let Q be as above, so that $T = \max(p_1, \dots, p_k)$. In light of the upper bound above, this is sufficient.

Example 6.3. Let p, q, r be distinct primes. Suppose that $A \in \mathcal{M}^+(\mathbb{Z}_M)$, where $pqr | M$, and that

$$(6.8) \quad \{pq, pr, qr\} \subset S_A.$$

By the definition of S_A , this implies that

$$(6.9) \quad (|A|, pqr) = 1,$$

and in particular none of Φ_p, Φ_q, Φ_r can divide A .

Let us try to apply the single-cluster construction in Section 3.1 in this case. With $\mathcal{C} := S_A$, at most one of p, q, r may divide Q , so that T must be divisible by at least two of them. Hence we need (at least) a bound of the form $|A| > \min(pq, pr, qr)$.

Corollary 6.5 does indeed provide such a bound. Applying the corollary to A with $m = q$, and using that Φ_{pq} and Φ_{qr} divide A but Φ_q does not, we see that $|A| \geq qr$. By (6.9), the inequality must be strict. Interchanging the primes, we get that

$$(6.10) \quad |A| > \max(pq, pr, qr).$$

This is in fact sufficient, with e.g., $Q = q$ and $T = pr < |A|$.

The following example shows that the lower bound in (6.10) is essentially optimal up to a multiplicative constant. Define $A \subset \mathbb{Z}_M$, where $pqr | M$, so that

$$A = (a * F_p * F_q) \cup (a' * F_q * F_r) \cup (a'' * F_p * F_r).$$

(With M/pqr sufficiently large, we can choose a, a', a'' so that they belong to different $D(M)$ -grids. Then the three ‘‘components’’ above are disjoint.) Then (6.8) holds, and $|A| = pq + pr + qr$, matching the order of magnitude of (6.10) if p, q, r have about the same size.

We note, however, that the argument does not extend to situations when two or more elements of S_A do not have a common ‘‘direct parent’’. For example, Corollary 6.5 does not provide any improvement of the single-step bound from Lemma 6.1 if $S_A = \{p^2q^3, p^4r^5, q^6r^7\}$, nor does the stronger version of it mentioned in Remark 6.1.

7. PROOF OF THEOREM 1.8

We will rely on the existing results on vanishing sums of roots of unity as in (1.1) with small k [28], [10]. A vanishing sum of the form (1.1) is called *minimal* if there is no proper subset $\{i_1, \dots, i_\ell\} \subsetneq \{1, \dots, k\}$ such that $z_{i_1} + \dots + z_{i_\ell} = 0$. For small k , all minimal vanishing sums of roots of unity with k elements can be classified and enumerated explicitly. Such an enumeration is provided in [28, Table 1] for $k \leq 12$, and extended in [10] to $k \leq 16$.

The results of [28], [10] can be converted to our language of arrays and fibers from Section 4 as follows. Recall from Section 1 that any vanishing sum of roots of unity can be phrased in terms of cyclotomic divisibility of polynomials. Conversely, if $A(X)$ is the mask polynomial corresponding to some $A \subset \mathbb{N}$ and $\Phi_N(X) | A(X)$ for some $N \in \mathbb{N} \setminus \{1\}$, then we must have

$$A(e^{2\pi i/N}) = \sum_{a \in A} e^{2\pi i a/N} = 0.$$

In other words, $A(e^{2\pi i/N})$ forms a vanishing sum of roots of unity. We may further reduce modulo N , so that $A \bmod N$ is a multiset in $\mathcal{M}^+(\mathbb{Z}_N)$. For each fixed N , the correspondence between multisets $A \bmod N$ in $\mathcal{M}^+(\mathbb{Z}_N)$ satisfying $\Phi_N(X) | A(X)$ and vanishing sums of roots of unity of the form

$$(7.1) \quad \sum_{a \in \mathbb{Z}_N} w(a) e^{2\pi i a/N} = 0$$

is one-to-one. The above sum is minimal if and only if A has the following minimality property:

Property (M): there is no multiset $A' \in \mathcal{M}^+(\mathbb{Z}_N)$ such that $A \neq A'$, $\Phi_N(X) | A'(X)$, and $w_{A'}(x) \leq w_A(x)$ for all $x \in \mathbb{Z}_N$.

If the sum is not minimal, its decomposition into minimal vanishing sums corresponds to writing $A(X)$ as a sum of polynomials $A_j(X)$ which have that property. This decomposition into minimal relations is the focus of [28] and [10]; that such a decomposition exists can be proven using induction upon the weight functions w appearing in the vanishing sums (7.1).

By the equivalence (i) \Leftrightarrow (iii) in Proposition 4.2, all polynomials $A(X)$ which do have the minimality property must fall into one of the following categories:

- $A(X)$ is an N -fiber in some direction.
- $A(X)$ is an “irreducible” linear combination of N -fibers as in Proposition 4.2 (iii) that cannot be expressed as a linear combination of N -fibers with nonnegative coefficients.

The first category of minimal sums is referred to in [28], [10] as R_p , where the prime p indicates the direction of the N -fiber. Such configurations can only occur when $p | N$. The second category is described using a language of recursive relations. For sets of cardinality at most 10, we will only need the $(R_p : nR_q)$ notation of [28]. Here, the integers p, q are distinct primes and we always have $1 \leq n < p$.

In our language, a given configuration of type $(R_p : nR_q)$ is constructed as follows. We choose some N with $2pq | N$ and work in \mathbb{Z}_N . Start with an N -fiber $x * F_p^N$ in the p direction. Choose n points x_1, \dots, x_n of that fiber, and subtract $x_j * F_q^N$ for each j . This “cancels” the points x_1, \dots, x_n , which now have weights 0, and introduces $n(q-1)$ points with negative weight -1 . Finally, add a N -fiber in the 2 direction through each point with weight -1 . This “cancels” all the negative weights, and introduces additional $n(q-1)$ points with weight 1. The total weight of the configuration, modulo \mathbb{Z}_N , is $(p-n) + n(q-1) = p + nq - 2n$.

Example 7.1. An example of a configuration of type $(R_5 : 2R_3)$ is provided by a multiset $\Xi \in \mathcal{M}^+(\mathbb{Z}_N)$ with the mask polynomial

$$(7.2) \quad \begin{aligned} \Xi(X) &= F_5^N(X) - (X^{N/5} + X^{2N/5})F_3^N(X) + \left(\sum_{i,j=1}^2 X^{\frac{iN}{5} + \frac{jN}{3}} \right) F_2^N(X) \\ &= 1 + X^{3N/5} + X^{4N/5} + \sum_{i,j=1}^2 X^{\frac{iN}{5} + \frac{jN}{3} + \frac{N}{2}}, \end{aligned}$$

where N is divisible by $2 \cdot 5 \cdot 3 = 30$. We have $\Xi(1) = 5 + 2 \cdot 3 - 2 \cdot 2 = 7$. Note that there are multiple configurations of the same type, depending on the placement of fibers. We illustrate one presentation below, which highlights the fiber-geometry underpinning equation (7.2).




FIGURE 1. We construct the above multiset $\Xi \in \mathcal{M}^+(\mathbb{Z}_{30})$ of cardinality 7 by combining 2-fibers, 3-fibers and 5-fibers with appropriately-chosen ± 1 weights.




FIGURE 2. Begin with any 30-fiber in the $p = 5$ direction. In this example, we chose $x * F_5^{30} = (0, 0, 2) * F_5^{30}$. Assign weight $+1$ to this fiber.




FIGURE 3. Select points x_{j_1}, x_{j_2} from the previous fiber and subtract $x_{j_1} * F_3^{30}$ and $x_{j_2} * F_3^{30}$. In this example, we chose $x_{j_1} = (1, 0, 2)$ and $x_{j_2} = (4, 0, 2)$. This introduces $n(q-1) = 2 \times 2 = 4$ points of weight -1 to our initial set.




FIGURE 4. We cancel the remaining points of the fibers $x_{j_1} * F_q^N$ and $x_{j_2} * F_q^N$ by adding $n(q-1)$ 30-fibers in the 2 direction with weight $+1$. After cancellation, this adds $n(q-1)$ points of weight $+1$ to our set.

We now turn to sets of small cardinality. By Proposition 2.1 and the proof of Theorem 1.7 in Section 2, it suffices to prove the following.

Lemma 7.2. *Let $A \subset \mathbb{N}$. Assume that $|A| \leq 10$, and that $S_A \neq \emptyset$. Then there exists a prime p_1 , relatively prime to $|A|$, such that:*

- (i) $p_1 | s$ for all $s \in S_A$,
- (ii) in the notation of Proposition 2.1, we have $p_1^{E_1} < |A|$.

Lemma 7.2 replaces Theorem 1.2 in the proof of Theorem 1.7. In this case, the cardinality of $|A|$ is given, and we just need to find an appropriate prime p_1 to use in Proposition 2.1. The lemma can fail for larger sets: for instance, Section 6.3 provides a counterexample with $|A| = 16$.

Proof of Lemma 7.2. Let A satisfy the assumptions of the lemma. Let $N \in S_A$ (note that N divides s_A but does not have to be equal to it). We write $A(X) \bmod X^N - 1$ as a sum of mask polynomials $A_j(X)$ satisfying the minimality condition (M) in \mathbb{Z}_N , each of which must be a fiber or an irreducible structure (as discussed previously). Since $N \in S_A$, we have $(N, |A|) = 1$. Therefore:

- If A_j is an N -fiber in the p direction for some $p|N$, then $(p, |A|) = 1$.
- If A_j is an irreducible structure of type $(R_p : nR_q)$ as described above, we must have $2pq|N$, hence $(2pq, |A|) = 1$.

Considering all the minimal structures listed in [28], and applying the above constraints, we are left with the following cases:

- $|A| = 5$, and A is a union of one N -fiber in the 2 direction and one N -fiber in the 3 direction,
- $|A| = 7$, and A is one of the following:
 - (a) a union of one N -fiber in the 2 direction and one N -fiber in the 5 direction,
 - (b) a union of two N -fibers in the 2 direction and one N -fiber in the 3 direction,
 - (c) an irreducible structure of type $(R_5 : 2R_3)$ (see Example 7.1),
- $|A| = 8$, and A is a union of one N -fiber in the 3 direction and one N -fiber in the 5 direction,
- $|A| = 9$, and A is a union of two N -fibers in the 2 direction and one N -fiber in the 5 direction,

- $|A| = 10$, and A is a union of one N -fiber in the 3 direction and one N -fiber in the 7 direction,

We first consider the case when A is a disjoint union of N -fibers in 2 different directions. Let $p_1 = 3$ if $|A| = 5$, $p_1 = 5$ if $|A| \in \{8, 9\}$, and $p_1 = 7$ if $|A| = 10$. We claim $E_1 = 1$, so that the conclusion of the lemma holds in each case.

Let α be the exponent of p_1 in the prime factorization of N ; since A contains an N -fiber in the p_1 direction, we must have $\alpha \geq 1$. Suppose that $s \in S_A$, $s \neq N$, and that the exponent of p_1 in the prime factorization of s is β . If $\beta > \alpha$, then $A \bmod s$ has nonempty intersection with at least p_1 different $D(s)$ -grids in \mathbb{Z}_s . By Lemma 4.4, we must have $\Phi_s | A \cap \Lambda$ for each such grid, hence $A \bmod s$ can be written as a union of at least p_1 non-empty multisets $A_j \in \mathcal{M}^+(\mathbb{Z}_s)$ with $\Phi_s | A_j$ for all j . But that is not compatible with the above list of permitted structures. If $\beta < \alpha$, the same argument applies, but with N and s interchanged.

Assume now that $|A| = 7$, and let $p_1 = 2$. For each $s \in S_A$, $A \bmod s$ has one of the structures listed in (a)-(c) above, hence $2|s$. To complete the proof of the lemma, we need to prove that $E_1 \leq 2$. We start with an auxiliary result.

Claim. Suppose that $B \subset \mathbb{N}$ satisfies $|B| \leq 4$, and that $\Phi_{s_1}(X)\Phi_{s_2}(X)|B(X)$ for some $s_1, s_2 \in \mathbb{N} \setminus \{1\}$ such that the exponents α_1, α_2 of 2 in the prime factorization of p_1 and p_2 satisfy $1 \leq \alpha_1 < \alpha_2$. Then one of the following holds:

- (i) $|B| = 3$, and $B \bmod s_i$ is a translate of $F_3^{s_i}$ for both $i = 1$ and $i = 2$,
- (ii) $|B| = 4$, and $B \bmod s_2$ is a translate of a ‘‘double fiber’’ in the 2 direction, with

$$(7.3) \quad B(X) = X^b(1 + X^{s_2/2})(1 + X^c) \pmod{X^{s_2} - 1}$$

for some $b, c \in \mathbb{Z}_{s_A}$ such that $c \equiv s_1/2 \pmod{s_1}$.

Proof of Claim. Since the smallest irreducible structure has 5 elements, $B \bmod s_i$ must be a union of s_i -fibers in the directions of 2 or 3. If $B \bmod s_1$ contains a translate of $F_3^{s_1}$, then $|B| = 3$ and we are in case (i). Suppose now that $B \bmod s_1$ is a union of (one or two) 2-fibers. Then $B \bmod s_1$ has nonempty intersection with at least two 2^{α_1} -grids in \mathbb{Z}_{s_1} . It follows that $B \bmod s_2$ has nonempty intersection with at least two 2^{α_1} -grids Λ_1, Λ_2 in \mathbb{Z}_{s_2} , and by Lemma 4.4, we must have $\Phi_{s_2} | B \cap \Lambda_i$ for $i = 1, 2$. Thus $B \cap \Lambda_i$ must be an s_2 -fiber in the 2 direction for each $i \in \{1, 2\}$. This implies the conclusion of case (ii). \square

We now return to the proof of Lemma 7.2 in the case $|A| = 7$, with $p_1 = 2$. Assume for contradiction that $s_i \in S_A$ for $i = 1, 2, 3$, with $1 \leq \alpha_1 < \alpha_2 < \alpha_3$, where α_i is the exponent of 2 in the prime factorization of A .

First, $A \bmod s_1$ must have one of the forms listed in (a)-(c) above with $N = s_1$. In each case, it follows that A has nonempty intersection with at least two 2^{α_1} -grids in \mathbb{Z}_{s_1} , therefore also in \mathbb{Z}_{s_2} . For each such grid Λ_j , $A \cap \Lambda_j$ must be divisible by $\Phi_{s_2}\Phi_{s_3}$, and at least one such set must have cardinality at most 3. Applying the claim above, we see that there can be at most two such grids Λ_1 and Λ_2 , and that (after possibly relabelling the grids) $A_i := A \cap \Lambda_i$ with $i = 1, 2$ satisfy, respectively, the conclusions (i) and (ii) of the claim.

We now return to \mathbb{Z}_{s_1} , where A_1 and A_2 are multisets with $|A_1| = 3$ and $|A_2| = 4$ contained in disjoint 2^{α_1} -grids. Considering the structures in (a)-(c), we see that only (c) is compatible with this.

Let γ_i be the exponent of 3 in the prime factorization of s_i for $i = 1, 2, 3$. Since A_1 must be a fiber in the 3 direction in both \mathbb{Z}_{s_2} and \mathbb{Z}_{s_3} , we have $\gamma_2 = \gamma_3$. Since the points of A_1 form a part of an s_1 -fiber in the 5 direction in \mathbb{Z}_{s_1} , we must have $\gamma_1 < \gamma_2$. Then, however, A_2 is contained in a single 3^{γ_1} -grid in \mathbb{Z}_{s_1} , so that A cannot form an $(R_5 : 2R_3)$ structure there. This exhausts all possible cases, and ends the proof of the lemma. \square

As mentioned in the introduction, it is likely that the cardinality bound $|A| \leq 10$ could be improved further with additional work along similar lines. However, for $|A| \geq 11$ the number of cases to consider increases rapidly, making the task significantly more time-consuming.

8. ACKNOWLEDGEMENT

The first author was supported by NSERC Discovery Grant 22R80520. The second author was supported by NSERC Discovery Grants 22R80520 and GR010263.

REFERENCES

- [1] M. Bateman, A. Volberg, *An estimate from below for the Buffon needle probability of the four-corner Cantor set*, *Math. Res. Lett.* 17 (2010), 959–967.
- [2] M. Bond, *Combinatorial and Fourier Analytic L^2 Methods For Buffon’s Needle Problem*, Ph, D, thesis, University of Michigan, <http://bondmatt.wordpress.com/2011/03/02/thesis-second-complete-draft/>.
- [3] M. Bond, I. Laba, and A. Volberg, *Buffon needle estimates for rational product Cantor sets*, *Amer. J. Math.* 136 (2014), 357–391.
- [4] M. Bond, A. Volberg: *Buffon needle lands in ϵ -neighborhood of a 1-dimensional Sierpinski Gasket with probability at most $|\log \epsilon|^{-c}$* , *Comptes Rendus Mathematique*, Volume 348, Issues 11–12, June 2010, 653–656.
- [5] M. Bond, A. Volberg: *Circular Favard Length of the Four-Corner Cantor Set*, *J. of Geometric Analysis* 21 (2011), 40–55.
- [6] M. Bond, A. Volberg, *Buffon’s needle landing near Besicovitch irregular self-similar sets*, *Indiana Univ. Math. J.* 61 (2012), no. 6, 2085–2109.
- [7] T. Bongers, *Geometric bounds for Favard length*, *Proc. Amer. Math. Soc.* 147, (2019)
- [8] T. Bongers, K. Taylor, *Transversal families of nonlinear projections and generalizations of Favard length*, available as preprint at arXiv:2105.01708
- [9] A. Chang, D. Dąbrowski, T. Orponen, M. Villa, *Structure of sets with nearly maximal Favard length*, preprint, arXiv:2203.01279
- [10] L. Christie, K. Dykema, I. Klep, *Classifying minimal vanishing sums of roots of unity*, preprint, arXiv:2008.11268
- [11] L. Cladek, B. Davey, K. Taylor, *Upper and lower bounds on the rate of decay of the Favard curve length for the four-corner Cantor set*, to appear in *Indiana U. Math. J.* (2022).
- [12] J.H. Conway, A.J. Jones: *Trigonometric diophantine equations (On vanishing sums of roots of unity)*, *Acta Arithmetica* 30 (1976), 229–240.
- [13] D. Coppersmith, J.P. Steinberger: *On the entry sum of cyclotomic arrays*, *Integers: the Electronic Journal of Combinatorial and Additive Number Theory*, 6 (2006), # A26.
- [14] N.G. de Bruijn, *On the factorization of cyclic groups*, *Indag. Math.* 15 (1953), 370–377.
- [15] R. Kenyon, *Projecting the one-dimensional Sierpinski gasket*, *Israel J. Math.* 79 (2006), 221–238.
- [16] G. Kiss, R. D. Malikiosis, G. Somlai, M. Vizer, *On the discrete Fuglede and Pompeiu problems*, *Analysis & PDE* 13 (2020), 765–788.
- [17] I. Laba, *Recent progress on Favard length estimates for planar Cantor sets*, in: *Operator-Related Function Theory and Time-Frequency Analysis*, Proceedings of the 2012 Abel Symposium, K. Grochenig, Y. Lyubarskii, K. Seip, eds., Springer 2015, pp. 117–145.
- [18] I. Laba, I. Londner, *Combinatorial and harmonic-analytic methods for integer tilings*, *Forum of Mathematics - Pi* 10:e8 (2022), 1–46.

- [19] I. Laba, K. Zhai: *The Favard length of product Cantor sets*, Bull. London Math. Soc. 42 (2010), 997–1009.
- [20] J.C. Lagarias, Y. Wang, *Tiling the line with translates of one tile*, Invent. Math. 124 (1996), 341–365.
- [21] J.C. Lagarias, Y. Wang, *Spectral sets and factorization of finite abelian groups*, J. Funct. Anal. 145 (1997), 73–98.
- [22] T.Y. Lam and K.H. Leung, *On vanishing sums of roots of unity*, J. Algebra 224 (2000), 91–109.
- [23] H. B. Mann, *On Linear Relations Between Roots of Unity*, Mathematika 12, Issue 2 (1965), 107–117.
- [24] P. Mattila, *Orthogonal projections, Riesz capacities, and Minkowski content*, Indiana Univ. Math. J. 124 (1990), 185–198.
- [25] P. Mattila: *Geometry of Sets and Measures in Euclidean Spaces*, Cambridge University Press, 1995.
- [26] F. Nazarov, Y. Peres, A. Volberg: *The power law for the Buffon needle probability of the four-corner Cantor set*, Algebra i Analiz 22 (2010), 82–97; translation in St. Petersburg Math. J. 22 (2011), 61–72.
- [27] Y. Peres, B. Solomyak, *How likely is buffon’s needle to fall near a planar Cantor set?*, Pacific J. Math. 24 (2002), 473–496.
- [28] B. Poonen and M. Rubinstein: *Number of Intersection Points Made by the Diagonals of a Regular Polygon*, SIAM J. Disc. Math. 11 (1998), 135–156.
- [29] L. Rédei, *Über das Kreisteilungspolynom*, Acta Math. Hungar. 5 (1954), 27–28.
- [30] L. Rédei, *Natürliche Basen des Kreisteilungskörpers*, Abh. Math. Sem. Univ. Hamburg 23 (1959), 180–200.
- [31] I. J. Schoenberg, *A note on the cyclotomic polynomial*, Mathematika 11 (1964), 131–136.
- [32] J. P. Steinberger, *Minimal vanishing sums of roots of unity with large coefficients*, Proc. London Math. Soc. 97 (3) (2008), 689–717.
- [33] T. Tao, *A quantitative version of the Besicovitch projection theorem via multiscale analysis*, Proc. London Math. Soc. 98 (2009), 559–584.
- [34] X. Tolsa, *Analytic capacity, rectifiability, and the Cauchy integral*, Proceedings of the ICM 2006, Madrid.
- [35] B. Wilson, *Sets with arbitrarily slow Favard decay*, preprint available at arxiv.org/abs/1707.08137

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA, VANCOUVER,
B.C. V6T 1Z2, CANADA

ilaba@math.ubc.ca, cmarshall@math.ubc.ca