

A SEMIDEFINITE FRAMEWORK FOR THE SIEVE

ZARATHUSTRA BRADY

ABSTRACT. We describe a semidefinite programming framework for proving upper bounds on concrete sifting problems, and show that the Large Sieve can be interpreted as a special case of this framework. With a small tweak, the Larger Sieve also falls into this framework.

We compare the semidefinite approach to the linear programming approach (i.e., the general framework of the combinatorial sieve and the Selberg sieve), and show that it has a qualitative advantage in a toy case where the primes are completely independent from each other. No new sieve-theoretic bounds are proved.

1. INTRODUCTION

Let Z be any finite set which we will apply the sieve to. Generally we will imagine that Z is a set of integers - in fact, most often Z will be a set of *consecutive* integers - but our framework applies to any finite set Z . We imagine that we are given a collection of partitions \mathcal{P} of the set Z . Generally, we imagine that \mathcal{P} corresponds to a set of primes $\{p_1, \dots\}$, and that the partition corresponding to the prime p_i partitions Z according to the congruence classes of the elements of Z modulo p_i . Our basic question is as follows.

Problem 1. Given a finite set Z and a set of partitions \mathcal{P} of Z , what upper bounds can we place on the size of a set $X \subseteq Z$ such that X misses at least one part of each partition $p \in \mathcal{P}$?

More generally, we can imagine that we have a number κ_p attached to each partition $p \in \mathcal{P}$, and we can ask for upper bounds on the sizes of sets $X \subseteq Z$ which avoid at least κ_p parts of p for each $p \in \mathcal{P}$.

As a motivating example of a problem that fits into our framework, we have the problem of finding large admissible tuples.

Definition 1. A set $H \subseteq \mathbb{Z}$ is *admissible* if for all primes p , H avoids at least one congruence class modulo p .

Taking $Z = \{0, \dots, N - 1\}$ and \mathcal{P} to be the set of partitions of Z into congruence classes modulo primes $p \leq N$, we see that finding the largest admissible tuple $H \subseteq \{0, \dots, N - 1\}$ is a special case of our general problem.

We approach this question from a somewhat unusual point of view. We imagine that we are handed a relatively small set Z and collection of partitions \mathcal{P} (i.e., $|Z|, |\mathcal{P}| < 10^{10}$), and wish to use a computer to automatically prove a numerical upper bound on the size of X . Finding the exact optimal upper bound by brute force is no good - we want an *efficient* algorithm for finding upper bounds. Failing that, we would be satisfied with proofs of upper bounds which can be efficiently and mechanically verified, but which may be very difficult to find. Ideally, proofs should be put into a standard form so that examples of proofs that work well can be compared to each other and, hopefully, generalized. This second requirement leads us to consider specialized proof frameworks which are less general than, say, all of axiomatic set theory.

This material is based upon work supported by the NSF Mathematical Sciences Postdoctoral Research Fellowship under Grant No. (DMS-1705177).

In Section 2 we will review the standard framework of systems of linear inequalities, together with a choice of *sieve weights* which are used to produce a proof of an upper bound on the size of a set $X \subseteq Z$ avoiding at least one congruence class from each partition in \mathcal{P} . This framework has several drawbacks - most importantly for us, it is not clear in general if one can efficiently verify that a collection of sieve weights truly leads to a valid inequality. In practice, this difficulty is avoided by making special choices of sieve weights, for which specialized arguments can be used to prove that they lead to valid inequalities.

After the review of the linear approach we will introduce a new framework of systems of quadratic inequalities in Section 3, which we will represent using *positive semidefinite matrices*. A proof in this framework is just a collection of matrices satisfying certain positive semidefiniteness conditions on their submatrices, such that their sum is a small multiple of the identity matrix. The entries of these matrices are analogous to the sieve weights of the linear framework.

The advantage of the semidefinite approach is that it is easy to mechanically check whether a matrix is positive semidefinite - for instance, one can simply compute the Cholesky decomposition of the matrix (in terms of the original quadratic inequalities, we can think of this as trying to prove a quadratic inequality by repeatedly completing the square). In fact, standard semidefinite programming algorithms allow one to compute the best possible upper bound using this framework in polynomial time (in numerical experiments, however, the standard semidefinite solvers have not been able to handle sets Z with $|Z| \gg 10^3$ in a reasonable amount of time).

In Sections 4 and 5, we will show that the Large Sieve and the Larger Sieve both fit into our general semidefinite framework, although the framework needs to be expanded slightly to handle the Larger Sieve. The special form of the semidefinite matrices used in these two cases leads to the consideration of a simpler (but slightly less powerful) semidefinite framework which is described in Section 6.

In Section 7, we show that there is a sense in which we can combine inequalities which each consider just a single partition on its own to produce strong upper bounds on the size of the sifted set X . This can be contrasted with the case of the linear framework, where considering one partition at a time only leads to a nontrivial upper bound if $\sum_{p \in \mathcal{P}} \frac{1}{|p|}$ is less than 1.

1.1. Notation. The author has attempted to find notation which reflects the general nature of the problem, but which matches with the usual multiplicative notation used in sieve theoretic arguments as closely as possible. We therefore always denote partitions in \mathcal{P} with lowercase letters p or q , and we describe subsets d of \mathcal{P} with a multiplicative notation. Thus 1 corresponds to the empty subset of \mathcal{P} , and we think of $d \cup k, d \cap k$ as corresponding to the lcm and gcd of d and k . If $d, k \subseteq \mathcal{P}$ are disjoint, then we write dk for their disjoint union. We also abuse notation by writing p for the singleton subset $\{p\} \subseteq \mathcal{P}$, so if $p \in \mathcal{P}, d \subseteq \mathcal{P}$ then the expression pd corresponds to the subset $d \cup \{p\}$ of \mathcal{P} .

For divisibility, the expression $p \in d$ corresponds to p being a prime dividing d . The expression $k \subseteq d$ corresponds to k being a divisor of d . The notation $d \setminus k$ for the set-theoretic difference is a curious anomaly: if $k \subseteq d$, then it corresponds to $\frac{d}{k}$, and in general it corresponds to $\frac{d}{\gcd(d,k)}$.

When working in the general setup, we use $|p|$ to denote the number of parts in the partition p . We extend κ to a multiplicative function on $d \subseteq \mathcal{P}$ by $\kappa(d) = \prod_{p \in d} \kappa_p$. We also define multiplicative functions $\phi(d), \mu(d)$ for $d \subseteq \mathcal{P}$ by $\phi(p) = |p| - 1$ and $\mu(p) = -1$ in analogy with the usual Euler ϕ -function and the Möbius function, although we will later need to introduce variants of the ϕ function which take the numbers κ_p as parameters.

A multiplicative notation for matrices will be introduced in Section 4, and this notation will prove useful in all later sections.

2. THE LINEAR FRAMEWORK (AND ITS DIFFICULTIES)

We will suppose throughout this section that the partitions in \mathcal{P} correspond to primes p , in order to make use of the existing multiplicative notation of divisibility and squarefree numbers, rather than introducing new notation for sets of partitions. Suppose that for each prime p , $X \subseteq Z$ avoids a set $A_p \subseteq Z$ corresponding to some congruence class modulo p . If we are working in the more general setting with a number κ_p of congruence classes to be avoided modulo each p , we assume that $A_p \subseteq Z$ corresponds to some collection of κ_p congruence classes modulo p .

We extend the notation A_p to squarefree numbers d multiplicatively:

$$A_d = \bigcap_{p|d} A_p,$$

where we interpret A_1 as Z . The linear framework assumes as given a system of linear inequalities

$$-R_d \leq |A_d| - \frac{\kappa(d)}{d}|Z| \leq R_d,$$

where R_d are *remainder terms* which are supplied to us by an outside source, and $\kappa(d)$ is defined by

$$\kappa(d) = \prod_{p|d} \kappa_p.$$

The reader should keep in mind that the main case of interest has $\kappa(d) = 1$ for all d .

Adding together the given inequalities with *sieve weights* λ_d , we get the inequality

$$\sum_d \lambda_d |A_d| = \left(\sum_d \lambda_d \frac{\kappa(d)}{d} \right) |Z| + \mathcal{O}^* \left(\sum_d |\lambda_d| R_d \right),$$

where the notation $f = \mathcal{O}^*(g)$ is a shorthand for the inequality $-g \leq f \leq g$. In order for this to give an upper bound on $|X|$, we need the following implication to be valid:

$$X \subseteq A_1, \forall p \ X \cap A_p = \emptyset \implies |X| \leq \sum_d \lambda_d |A_d|.$$

This implication is equivalent to $\lambda_1 \geq 1$ together with the system of inequalities

$$(*) \quad \forall k \quad \sum_{d|k} \lambda_d \geq 0.$$

So the general linear framework proceeds as follows. First we estimate the remainder terms R_d . Then we choose a collection of sieve weights λ_d which satisfy $(*)$ and have $\lambda_1 = 1$. Finally, we deduce the upper bound

$$|X| \leq \left(\sum_d \lambda_d \frac{\kappa(d)}{d} \right) |Z| + \sum_d |\lambda_d| R_d.$$

The biggest difficulty with the linear framework is that we need to verify the system of inequalities $(*)$. Even if the λ_d are supported on values of d with $d \ll |Z|$, we still have to check $(*)$ for *all* k dividing $\prod_{p \leq |Z|} p$.

Proposition 1. *It is co-NP-hard to check whether a given collection of sieve weights λ_d satisfies $(*)$, even if λ_d is supported on d with at most two prime factors.*

Proof. In fact, this problem is equivalent to determining whether a quadratic polynomial ever takes a negative value, when its inputs are restricted to values from $\{0, 1\}$: if we write $k = \prod_p p^{x_p}$, then $\sum_{d|k} \lambda_d = \sum_{p < q} \lambda_{pq} x_p x_q + \sum_p \lambda_p x_p + \lambda_1$ (note that $x_p^2 = x_p$ for $x_p \in \{0, 1\}$, so there is no loss of generality in considering quadratic polynomials of this form). This problem is known as the Binary Quadratic Programming problem, and is well-known to be hard. For instance, we can reduce from

k -clique as follows: let G be a graph with vertices labeled by primes p , put $\lambda_1 = 1$, $\lambda_p = -\frac{1}{k-1}$, and $\lambda_{pq} = |G|$ for any p, q which are not connected by an edge of the graph G . \square

For computing asymptotics, a good strategy is to bucket the large primes p into finitely many buckets B_i , such that each bucket has $\sum_{p \in B_i} \frac{1}{p}$ small (the small primes can be handled by a Selberg sieve, and do not contribute much to the asymptotics). Then we decide on values for λ_d which only depend on the number of primes from each bucket which divide d , so we can write

$$\lambda_d = \lambda(e_1, \dots, e_k), \quad e_i = \#\{p \in B_i \text{ s.t. } p \mid d\}.$$

The system of inequalities (*) then becomes

$$\forall n_1, \dots, n_k \in \mathbb{N} \quad \sum_{e_1, \dots, e_k} \lambda(e_1, \dots, e_k) \binom{n_1}{e_1} \cdots \binom{n_k}{e_k} \geq 0,$$

that is, we must determine whether a given polynomial in k variables is nonnegative when its variables take values from the natural numbers.

Proposition 2. *If k is sufficiently large, then there is no algorithm which determines whether a given polynomial $f \in \mathbb{Z}[x_1, \dots, x_k]$ is nonnegative on the naturals, even if the degree of f is bounded by 8.*

Proof. By Matiyasevich's resolution of Hilbert's Tenth Problem [4], if k is sufficiently large then there is no algorithm which determines whether a given polynomial $g \in \mathbb{Z}[x_1, \dots, x_k]$ ever takes the value 0 for natural inputs x_1, \dots, x_k , even if the degree of g is bounded by 4. Now take $f(x_1, \dots, x_k) = g(x_1, \dots, x_k)^2 - 1$, and note that f is nonnegative on the naturals if and only if g never takes the value 0 for natural inputs. \square

Despite the above result, it turns out that by slightly increasing the high-order terms of our polynomial $\sum_{e_1, \dots, e_k} \lambda(e_1, \dots, e_k) \binom{n_1}{e_1} \cdots \binom{n_k}{e_k}$ we can reduce proving the inequality to a finite search together with a proof that the polynomial goes to infinity as the variables go to infinity, at the cost of slightly decreasing the quality of our asymptotic bounds. This enabled Selberg [9] to prove that optimal asymptotics for the linear approach to the sieve are at least *computable* in theory.

In practice, rather than compute the optimal choice of sieve weights, most work on the linear approach to the sieve focuses on sets of sieve weights which are guaranteed to work by some simple principle. The most prominent example is the principle behind the Selberg sieve: at a high-level, rather than trying to optimize over the set of polynomials which are nonnegative on the natural numbers, we try to optimize over the set of polynomials which can be written as a sum of squares (and are thus nonnegative on the reals as well). The alternative approach is to use recursive principles such as Buchstab iteration, which are guaranteed to produce new valid sets of sieve weights from old valid sets of sieve weights.

There are a few additional drawbacks to the linear framework from the point of view of the general problem considered in this paper. The first is that the requirement of bounds of the form

$$|A_d| = \frac{\kappa(d)}{d} |Z| + \mathcal{O}^*(R_d)$$

is very inflexible. In cases where the parts of some partition $p \in \mathcal{P}$ have very different sizes from each other, the bounds we get fail to degrade gracefully. A more detailed linear relaxation (with different sieve weights for each element of Z) could overcome this difficulty.

The next drawback we will mention is related to the concept of the *hierarchy* of successively stronger linear relaxations to a problem. The *basic* LP would only allow us to use sieve weights λ_d with d equal to a prime. At the second level of the hierarchy, we would consider sieves with λ_d supported on d a product of two primes, and so on. A sieve which is low on this hierarchy is

considered logically simpler than a sieve which is high on this hierarchy. Unfortunately, in the case of ordinary sieving, any constant level of this hierarchy of linear relaxations is useless.

Proposition 3. *If $\sum_p \frac{\kappa_p}{p}$ diverges, then for any set of sieve weights λ_d satisfying (*) (and with $\lambda_1 = 1$) which are supported on the set of d with at most $2m$ prime factors, we have*

$$\frac{\sum_d \lambda_d \frac{\kappa(d)}{d}}{\prod_p \left(1 - \frac{\kappa_p}{p}\right)} \rightarrow \infty.$$

Proof. Letting $\theta(k) = \sum_{d|k} \lambda_d$, we have

$$\prod_p \left(1 - \frac{\kappa_p}{p}\right) \sum_k \theta(k) \frac{\kappa(k)}{k} = \sum_d \lambda_d \frac{\kappa(d)}{d}$$

by Möbius inversion. So we just have to show that $\sum_k \theta(k) \frac{\kappa(k)}{k} \rightarrow \infty$. By (*), each $\theta(k)$ is ≥ 0 , so we just have to show that enough of them are sufficiently large to finish.

For any k having $2m + 1$ prime factors, we have

$$0 = \lambda_k = \sum_{d|k} \mu(d) \theta(d)$$

by Möbius inversion again, so

$$\sum_{d|k, \mu(d)=-1} \theta(d) = \sum_{d|k, \mu(d)=1} \theta(d) \geq \theta(1) = 1.$$

Thus, for every k with $2m + 1$ prime factors, k has some nontrivial divisor d with $\theta(d) \geq 2^{-2m}$. Call d “good” if $\theta(d) \geq 2^{-2m}$, otherwise call it “bad”.

Suppose for contradiction that $\sum_k \theta(k) \frac{\kappa(k)}{k}$ remains bounded. Then in particular $\sum_{d \text{ good}} \frac{\kappa(d)}{d}$ remains bounded, so there is some bounded prime p_1 which is bad. Call d “ p_1 -good” if $p_1 \nmid d$ and one of $d, p_1 d$ is good. Then $\sum_{d \text{ } p_1\text{-good}} \frac{\kappa(d)}{d}$ remains bounded, so there is a bounded p_1 -bad prime p_2 . Continuing like this, we see that there is some bounded k with $2m + 1$ prime factors such that every nontrivial divisor of k is bad, a contradiction. \square

As a consequence, any set of sieve weights which work well for the ordinary sifting scenario is necessarily somewhat intricate, with the support of the sieve weights depending on the sizes of the prime factors as well as their number.

Finally, there is the famous *parity problem* identified by Selberg [9] in the case where all the κ_p are 1. The analysis goes roughly as follows. First, by analyzing the identity

$$\prod_p \left(1 - \frac{1}{p}\right) \sum_k \frac{\theta(k)}{k} = \sum_d \frac{\lambda_d}{d}$$

used in the proof of the previous proposition, we see that if the ratio $\frac{\sum_d \lambda_d/d}{\prod_p (1-1/p)}$ is reasonably small, then the average values of the $\theta(k)$ s (weighted by $\frac{1}{k}$) must also be small. By the Möbius inversion formula, the $|\lambda_d|$ s must be small on average as well.

So long as the support of the λ_d s is reasonable, this means that the contribution of the remainder terms in our bound

$$|X| \leq \left(\sum_d \frac{\lambda_d}{d}\right) |Z| + \sum_d |\lambda_d| R_d.$$

is small compared to the main term even if the remainders get somewhat large, especially if we consider using the same collection of sieve weights on a slightly larger set Z .

The fact that the bounds do not depend much on the remainder terms means that the linear framework can't make very effective use of tight bounds on the remainder terms R_d - bounds which are $O(1)$ do not lead to substantially better asymptotics than bounds which are $O((|Z|/d)^{1-\epsilon})$. So we see that it is possible to perturb the set Z by removing all numbers from Z which have an even number of prime factors and double counting the rest, without changing the fact that the remainder terms R_d are $O((|Z|/d)^{1-\epsilon})$. Since the upper bound applies to the perturbed case as well, the best possible upper bounds in this framework are necessarily a factor of 2 larger than the true bound if the support of the sieve weights is reasonable, even if the remainders are completely nonexistent.

3. THE SEMIDEFINITE FRAMEWORK

The main idea is to try to prove an inequality of the form

$$\left(\sum_{i \in X} x_i\right)^2 \leq \lambda \sum_{i \in X} x_i^2,$$

where X , the support of variables x_i , avoids at least one part of each partition P_i . Taking $x_i = 1$ for $i \in X$, this will prove the inequality $|X| \leq K$.

In order to prove an inequality of this form, we try to write the expression $\lambda \sum_i x_i^2 - (\sum_i x_i)^2$ as a sum of "obvious" inequalities $Q^d(\bar{x}) \geq 0$, where the quadratic form Q^d has the property that expressions of the form

$$Q^d(x_1, x_2, 0, x_4, \dots)$$

are positive semidefinite, where 0s have been inserted into indices which fall into the parts of the partitions corresponding to d which we are avoiding (recall that we are using a multiplicative notation, so d corresponds to a collection of partitions), for every way of choosing one part of each partition from d to avoid. When dealing with the Larger Sieve, it will also be helpful to include an additional inequality of the form $Q^0(\bar{x}) \geq 0$, where the only assumption on Q^0 is that its coefficients are nonnegative.

When manipulating quadratic forms on a computer, it is natural to organize their coefficients into symmetric matrices. To each quadratic form Q^d we associate a symmetric matrix B^d - we think of the system of matrices B^d as an analogue of a system of sieve weights λ_d from the linear framework. Let J be the all-ones matrix. Our proof of an upper bound on X will now take the form of a chain of inequalities:

$$x^T J x \leq \sum_d x^T B^d x + x^T J x = \lambda x^T x,$$

where each $x^T B^d x \geq 0$ can be easily checked in isolation.

Note that for any B, x we have $x^T B x = \text{Tr}(B x x^T)$, so it is natural to write $A = x x^T$ and rewrite our proof in the form

$$\text{Tr}(J A) \leq \sum_d \text{Tr}(B^d A) + \text{Tr}(J A) = \lambda \text{Tr}(A).$$

In this more general form, the positive semidefinite matrix A has a new interpretation. We can imagine having a probability distribution μ on the collection of all subsets $X \subseteq Z$, and let A be the covariance matrix:

$$A_{ij} = \mathbb{P}_\mu[i \in X \wedge j \in X].$$

To capture the fact that the unknown set X must avoid at least one part of each partition, we require that for every small $d \subseteq \mathcal{P}$, the matrix A is in the convex hull of matrices of the form $x x^T$, where the vector x is supported on a set which avoids at least one part of each partition $P \in d$.

An important special case to think about is the case where we just have one partition, and each part of the partition has size 1. For the sake of concreteness, imagine that our universe Z is \mathbb{Z}/p for some p . We define two closed convex cones:

$$\begin{aligned}\mathcal{A}_p &= \text{Conv}\{xx^T \text{ s.t. } x \in \mathbb{R}^p, x_i = 0 \text{ for some } i \in \mathbb{Z}/p\}, \\ \mathcal{B}_p &= \{B \in \mathbb{R}^{p \times p} \text{ s.t. } B|_{Z \setminus \{i\}} \succeq 0 \text{ for all } i \in \mathbb{Z}/p\},\end{aligned}$$

where $B|_{Z \setminus \{i\}}$ is a submatrix of B obtained by deleting the i th row and column of B , and $M \succeq 0$ means that the matrix M is positive semidefinite.

Proposition 4. *The cones $\mathcal{A}_p, \mathcal{B}_p$ defined above are dual:*

$$\begin{aligned}B \in \mathcal{B}_p &\iff \forall A \in \mathcal{A}_p, \text{Tr}(AB) \geq 0, \\ A \in \mathcal{A}_p &\iff \forall B \in \mathcal{B}_p, \text{Tr}(AB) \geq 0.\end{aligned}$$

Furthermore, for any matrices A, B we can efficiently check whether $B \in \mathcal{B}_p$ or $A \in \mathcal{A}_p$. For $p \geq 3$, both \mathcal{A}_p and \mathcal{B}_p contain a neighborhood of the identity matrix.

Proof. A matrix B is in \mathcal{B}_p iff for every i and every $x \in \mathbb{R}^p$ with $x_i = 0$, we have $x^T B x \geq 0$, and from $x^T B x = \text{Tr}(x x^T B)$ we get the first duality. We can check whether $B \in \mathcal{B}_p$ by computing Cholesky decompositions of the matrices $B|_{Z \setminus \{i\}}$ for each $i \in \mathbb{Z}/p$, this takes time roughly $O(p \cdot p^3) = O(p^4)$.

The statements about \mathcal{A}_p can be deduced from the statements about \mathcal{B}_p by conic duality, so long as we can show that \mathcal{A}_p is closed. Concretely, a matrix A is in \mathcal{A}_p iff there are matrices X_1, \dots, X_p such that

- $X_i \succeq 0$ for all i ,
- the i th row and column of X_i is all 0s, and
- $A = \sum_i X_i$.

Since each $X_i \preceq A$, if A is from a bounded set then the possible X_i s which we might consider also come from a bounded set, and this together with the fact that $X_i \succeq 0$ and $\sum_i X_i = A$ are closed conditions imply that \mathcal{A}_p is closed.

We can check whether a matrix $A \in \mathcal{A}_p$ by solving the semidefinite program corresponding to the three bullet points above. Alternatively, we can check $A \in \mathcal{A}_p$ by minimizing $\text{Tr}(AB)$ over $B \in \mathcal{B}_p$ (perhaps with some extra constraint such as $\|B\| \leq 1$).

That \mathcal{B}_p contains a neighborhood of the identity matrix is clear from the definition. For \mathcal{A}_p , we use the fact that the set of X_i as in the bullet points forms a neighborhood of $\frac{1}{p-1}I|_{Z \setminus \{i\}}$ in $\mathbb{R}^{(p-1) \times (p-1)}$, and the sum of these neighborhoods forms a neighborhood of I for $p \geq 3$. \square

Returning to the general case, we can define cones $\mathcal{A}_d, \mathcal{B}_d$ for any set d of partitions from \mathcal{P} .

Definition 2. For $d = \{p_1, \dots, p_k\} \subseteq \mathcal{P}$, where \mathcal{P} is a set of partitions of Z , we define the cones $\mathcal{A}_d, \mathcal{B}_d$ by

$$\begin{aligned}\mathcal{A}_d &= \text{Conv}\{xx^T \text{ s.t. } x \in \mathbb{R}^Z, \forall p_i \in d \exists c_i \in p_i \text{ with } x_j = 0 \text{ for } j \in c_i\}, \\ \mathcal{B}_d &= \{B \in \mathbb{R}^{Z \times Z} \text{ s.t. } \forall c_1 \in p_1, \dots, c_k \in p_k \ B|_{Z \setminus (c_1 \cup \dots \cup c_k)} \succeq 0\},\end{aligned}$$

where $B|_{Z \setminus (c_1 \cup \dots \cup c_k)}$ is the matrix formed by deleting all rows and columns in $c_1 \cup \dots \cup c_k$ from B .

A similar argument to the previous proposition shows that $\mathcal{A}_d, \mathcal{B}_d$ are dual for each d . The general result we will use is the following application of semidefinite duality.

Theorem 1. *If $D = \{d_1, \dots\}$ is any collection of subsets of \mathcal{P} , \mathcal{P} a set of partitions of Z , then*

$$(*) \quad \max_{A \in \bigcap_{d \in D} \mathcal{A}_d} \frac{\text{Tr}(AJ)}{\text{Tr}(A)} = \min_{B^{d_i} \in \mathcal{B}_{d_i}} \left\| \sum_{d \in D} B^d + J \right\|_{op}.$$

The minimum is attained with a collection of matrices B^{d_i} such that $\sum_{d \in D} B^d + J$ is a multiple of the identity matrix.

Proof. First we show that the left hand side is at most the right hand side. If $A \in \bigcap_{d \in D} \mathcal{A}_d$, then for each $B^{d_i} \in \mathcal{B}_{d_i}$ we have $\text{Tr}(AB^{d_i}) \geq 0$, so

$$\text{Tr}(AJ) \leq \sum_{d \in D} \text{Tr}(AB^d) + \text{Tr}(AJ) = \text{Tr} \left(\left(\sum_{d \in D} B^d + J \right) A \right) \leq \left\| \sum_{d \in D} B^d + J \right\|_{op} \text{Tr}(A),$$

where the last inequality follows from the fact that $|\text{Tr}(MA)| \leq \|M\|_{op} \text{Tr}(A)$ for any square matrix M and any positive semidefinite matrix A (to see this, just represent A as a sum of outer products xx^T).

For the other direction, we will use the fact that the cone $\bigcap_{d \in D} \mathcal{A}_d$ is dual to the cone $\sum_{d \in D} \mathcal{B}_d$, which will follow from the fact that each pair $\mathcal{A}_{d_i}, \mathcal{B}_{d_i}$ are dual as long as we can show that $\sum_{d \in D} \mathcal{B}_d$ is closed. In the case where no partition $p \in \mathcal{P}$ has just two parts, this follows from the facts that each $B^{d_i} \in \mathcal{B}_{d_i}$ has nonnegative entries along the diagonal and that the off-diagonal entries are bounded in terms of the diagonal entries. If some $p \in \mathcal{P}$ has just two parts, then some of the entries of matrices $B^{d_i} \in \mathcal{B}_{d_i}$ can take arbitrary values independently of the other entries, so these entries may be ignored, and all other entries of B^{d_i} are again bounded in terms of the diagonal entries.

Note that since $\text{Tr}(A) \geq 0$ for $A \in \bigcap_{d \in D} \mathcal{A}_d$, the maximum $\max_{A \in \bigcap_{d \in D} \mathcal{A}_d} \frac{\text{Tr}(AJ)}{\text{Tr}(A)}$ is the same as

$$\max_{\substack{A \in \bigcap_{d \in D} \mathcal{A}_d, \\ \text{Tr}(A)=1}} \text{Tr}(AJ),$$

and since the set $\{A \in \bigcap_{d \in D} \mathcal{A}_d, \text{Tr}(A) = 1\}$ is compact and nonempty, the maximum is actually attained. By the duality between $\bigcap_{d \in D} \mathcal{A}_d$ and $\sum_{d \in D} \mathcal{B}_d$, the maximum above is equal to the minimum $\lambda \in \mathbb{R}$ such that $\lambda I - J \in \sum_{d \in D} \mathcal{B}_d$ (and this minimum is attained because $\sum_{d \in D} \mathcal{B}_d$ is closed). For this minimum value of λ , there exist B^{d_i} with

$$\sum_{d \in D} B^d + J = \lambda I. \quad \square$$

Corollary 1. *If some $d \in D$ is equal to all of \mathcal{P} , then (*) equal to the size of the largest set $X \subseteq Z$ which avoids at least one part of each partition $p \in \mathcal{P}$.*

Proof. We just have to check that for any $A \in \mathcal{A}_{\mathcal{P}}$, we have $\text{Tr}(AJ) \leq \lambda \text{Tr}(A)$, where λ is the size of the largest such set X . Since each $A \in \mathcal{A}_{\mathcal{P}}$ can be written as a sum of rank one matrices xx^T with x supported on a set X which avoids at least one part of each partition $p \in \mathcal{P}$, we just need to show that if x is supported on X then $x^T J x \leq |X| x^T x$, that is, that $(\sum_{i \in X} x_i)^2 \leq |X| \sum_{i \in X} x_i^2$. This follows from Cauchy-Schwarz. \square

Of course, it is impractical to compute the sets $\mathcal{A}_{\mathcal{P}}, \mathcal{B}_{\mathcal{P}}$ if \mathcal{P} is large. So our strategy is to pick some collection $D = \{d_1, \dots\}$ of *small* subsets of \mathcal{P} , and hope that the intersection $\bigcap_{d \in D} \mathcal{A}_d$ is a good approximation to $\mathcal{A}_{\mathcal{P}}$. Then (*) can be efficiently computed to any desired accuracy (in theory), and sets of witnessing matrices B^{d_i} can be produced to give a proof of an upper bound.

Proposition 5. *If D is any collection of subsets of \mathcal{P} , then (*) can be computed to any desired accuracy ϵ in time polynomial in $\log(1/\epsilon)$, $|Z|$, and $\sum_{d \in D} \prod_{p \in d} |p|$.*

Proof. This follows from a general result of Khachiyan [3], which reduces the task of minimizing a convex function over a closed convex set to the task of testing whether a point is in the set. To test whether a matrix $B^d \in \mathcal{B}_d$, we just need to compute a Cholesky decomposition of $\prod_{p \in d} |p|$ submatrices of B . \square

All of the above results and definitions can be extended to the case where we sieve out κ_p parts of the partition p instead of just sieving out one part, but there is an annoying additional complication if κ_p is not bounded. To understand this complication, we again reduce to a simple case where we have just one partition p , where each part of p has just one element. We define closed convex cones

$$\mathcal{A}_{p,\kappa} = \text{Conv}\{xx^T \text{ s.t. } x \in \mathbb{R}^p, x_{i_1} = \dots = x_{i_\kappa} = 0 \text{ for some distinct } i_1, \dots, i_\kappa \in p\},$$

$$\mathcal{B}_{p,\kappa} = \{B \in \mathbb{R}^{p \times p} \text{ s.t. } B|_{Z \setminus \{i_1, \dots, i_\kappa\}} \succeq 0 \text{ for all distinct } i_1, \dots, i_\kappa \in p\},$$

and note that these are again dual to each other.

Proposition 6. *If p and κ are allowed to have the same order of magnitude, then it is NP-hard to test whether a given matrix B is contained in $\mathcal{B}_{p,\kappa}$.*

Proof. Checking whether $B \in \mathcal{B}_{p,\kappa}$ is the same as determining whether there exists $x \in \mathbb{R}^p$ having at most $p - \kappa$ nonzero entries and satisfying $x^T B x < 0$. If we could solve this efficiently for all κ , then we could find the sparsest vector x which satisfies $x^T B x < 0$. We will show this is NP-hard by a reduction from the problem of finding the sparsest vector that approximately solves a linear system, aka the “best subset selection problem” (see [7]).

The reduction from the best subset selection problem to our problem goes as follows: if it is hard to find a sparse x satisfying $\|Ax - b\| < 1$, then it is equally hard to find a sparse vector (x_0, x_1, \dots, x_n) satisfying $\|A(x_1, \dots, x_n)^T - bx_0\|^2 - x_0^2 < 0$, and the left hand side is a quadratic form in x_0, \dots, x_n . \square

In practice, when κ_p is allowed to be large, we get around this difficulty by restricting attention to special matrices B^p which have a symmetry property that allows us to efficiently verify that $B^p \in \mathcal{B}_{p,\kappa_p}$.

4. THE LARGE SIEVE

To write down the Large Sieve compactly, it is efficient to use a multiplicative notation for matrices. We think of our interval Z as a subset of $\prod_p \mathbb{Z}/|p|$. We use I_p, J_p for the $p \times p$ identity matrix and matrix of all ones, respectively, and if $M_p \in \mathbb{R}^{p \times p}, N_q \in \mathbb{R}^{q \times q}$, then we write $M_p \otimes N_q$ for their tensor product, thought of as a matrix with rows and columns indexed by elements of $\mathbb{Z}/|p| \times \mathbb{Z}/|q|$. In particular, we have

$$I_p \otimes I_q = I_{pq}, \quad J_p \otimes J_q = J_{pq}.$$

The matrices we write down with this multiplicative notation will, strictly speaking, have rows and columns indexed by elements of $\prod_p \mathbb{Z}/p$ which lie outside the interval Z . We ignore these extra rows and columns - that is, we consider the natural projection map $\mathbb{R}^{\prod_p p} \rightarrow \mathbb{R}^Z$ and the inclusion $\mathbb{R}^Z \rightarrow \mathbb{R}^{\prod_p p}$ corresponding to $Z \rightarrow \prod_p \mathbb{Z}/p$, and pre- and post-compose with these to obtain matrices in $\mathbb{R}^{Z \times Z}$.

Proposition 7. *If A, B, C, D are symmetric matrices with $A \succeq B \succeq 0$ and $C \succeq D \succeq 0$, then $A \otimes C \succeq B \otimes D \succeq 0$.*

Proof. To see that the tensor product of positive semidefinite matrices is positive semidefinite, just note that every positive semidefinite matrix is a positive combination of rank one matrices of the form xx^T , and that tensor products of such matrices are again of that form. To finish, note that

$$A \otimes C - B \otimes D = A \otimes (C - D) + (A - B) \otimes D \succeq 0. \quad \square$$

Lemma 1 (Montgomery [5]). *If $A \in \mathcal{A}_{d,\kappa}$, then*

$$\left(\prod_{p \in d} \frac{\kappa_p}{|p| - \kappa_p} \right) \text{Tr}(AJ) \leq \text{Tr} \left(A \left(\bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q \right) \right).$$

In other words, the matrix $B^d = \left(\bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q \right) - \left(\prod_{p \in d} \frac{\kappa_p}{|p| - \kappa_p} \right) J$ is in $\mathcal{B}_{d, \kappa}$.

Proof. Due to the multiplicative nature of the inequality, it's enough to check it in the case $d = \{p\}$ for just one partition p . So we just need to check that $|p|I_p - J_p - \frac{\kappa_p}{|p| - \kappa_p} J_p$ is in $\mathcal{B}_{p, \kappa_p}$.

Multiplying through by $\frac{|p| - \kappa_p}{|p|}$, this is equivalent to $(|p| - \kappa_p)I_p - J_p \in \mathcal{B}_{p, \kappa_p}$. If we delete any κ_p rows and corresponding columns of this matrix, then due to the symmetry of I_p, J_p the resulting matrix will be $(|p| - \kappa_p)I_{|p| - \kappa_p} - J_{|p| - \kappa_p}$, which is positive semidefinite by Cauchy-Schwarz. \square

Lemma 2 (Montgomery [5]). *If $S(\alpha)$ is the symmetric matrix such that*

$$x^T S(\alpha) x = \left| \sum_j e^{2\pi i j \alpha} x_j \right|^2,$$

that is, if $S(\alpha)_{jk} = \cos(2\pi \alpha(j - k))$, then

$$\bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q = \sum_{a \in (\mathbb{Z}/|d|)^\times} S\left(\frac{a}{|d|}\right).$$

Theorem 2 (Analytic Large Sieve Inequality [6]). *If $S(\alpha) \in \mathbb{R}^{Z \times Z}$ is defined as in the previous proposition, the α s are δ -spaced, and Z is an interval with $|Z| = N$, then*

$$\left\| \sum_{\alpha} S(\alpha) \right\|_{op} \leq N + \delta^{-1} - 1.$$

Corollary 2. *The matrix $B^1 = (N + \delta^{-1} - 1)I - \sum_{d \in D} \bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q$ is positive semidefinite if the fractions with denominators $|d|$ for $d \in D$ are δ -spaced. If the maximum $|d|$ for $d \in D$ is Q , then we can take $\delta^{-1} = Q(Q - 1)$.*

Putting it all together, and scaling everything down by a factor of $\sum_{d \in D} \frac{\kappa(d)}{\phi_\kappa(d)}$, where ϕ_κ is the multiplicative function with $\phi_\kappa(p) = |p| - \kappa_p$, we get the following proof of the (sieve-theoretic) Large Sieve inequality:

$$\begin{aligned} & \frac{1}{\sum_d \frac{\kappa(d)}{\phi_\kappa(d)}} \left(\left((N + \delta^{-1} - 1)I - \sum_{d \in D} \bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q \right) \right. \\ & \quad \left. + \sum_{d > 1} \left(\left(\bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q \right) - \frac{\kappa(d)}{\phi_\kappa(d)} J \right) \right) + J = \frac{N + \delta^{-1} - 1}{\sum_d \frac{\kappa(d)}{\phi_\kappa(d)}} I. \end{aligned}$$

Note that from our computational point of view, the analytic large sieve inequality (which is the hardest part of the proof) is considered a trivality, since in any particular case we may easily use a computer to numerically verify that the matrix

$$B^1 = (N + \delta^{-1} - 1)I - \sum_{\alpha} S(\alpha) = (N + \delta^{-1} - 1)I - \sum_{d \in D} \bigotimes_{p \in d} (|p|I_p - J_p) \bigotimes_{q \notin d} J_q$$

is positive semidefinite. In fact, the final set of matrices B^d used makes no explicit mention of the matrices $S(\alpha)$ whatsoever - they were only introduced in order to facilitate the proof that B^1 is positive semidefinite.

5. THE LARGER SIEVE

To treat Gallagher's Larger Sieve [1] in this framework, it is necessary to introduce an additional matrix B^0 where the only constraint on the entries of B^0 is that they are nonnegative. This corresponds to the fact that for the matrices $A = xx^T$ of interest to us, the vector x and the matrix

A will actually have positive entries (in fact, we are really only concerned with the case where the vector x has all of its entries in $\{0, 1\}$, but there doesn't seem to be any good way to make use of that fact). Here is Gallagher's sieve:

$$B^0 + \sum_p \frac{\left(\log(p)I_p - \frac{\log(p)}{p-\kappa_p}J_p\right) \otimes_{q \neq p} J_q}{\sum_q \frac{\log(q)}{q-\kappa_q} - \log(N)} + J = \frac{\sum_p \log(p) - \log(N)}{\sum_p \frac{\log(p)}{p-\kappa_p} - \log(N)} I.$$

The diagonal entries of B^0 are 0, while the off-diagonal entries are given by

$$B_{ij}^0 = \frac{\log(N) - \sum_{p|i-j} \log(p)}{\sum_p \frac{\log(p)}{p-\kappa_p} - \log(N)} \text{ for } i \neq j.$$

Note that the matrix B^0 is quite far from being positive semidefinite - every 2×2 submatrix with corresponding rows and columns has the form $\begin{bmatrix} 0 & x \\ x & 0 \end{bmatrix}$ for some $x \geq 0$.

In the case where $\kappa_p \approx \alpha p$ for some constant $0 \leq \alpha < 1$, if we apply the Larger Sieve with the first $eN^{1-\alpha}$ primes, we get an upper bound of approximately $e(1-\alpha)N^{1-\alpha} \log(N^{1-\alpha})$ on the size of the sifted set X . In the special case $\kappa_p = 1$ of interest to us, we get the rather weak upper bound $eN \log(N)$ on the size of the subset X of our interval Z (which has size only N).

6. A SIMPLER SEMIDEFINITE FRAMEWORK

Note that in both the Large Sieve and the Larger Sieve, the matrices B^d were highly symmetric - in the sense that they could be represented as linear combinations of tensor products of I and J - for all d other than 0, 1. Thus it seems likely that restricting our choices of B^d to highly symmetric matrices for $d \neq 0, 1$ does not lose too much power.

In this section we will show that testing whether $B^d \in \mathcal{B}_d$ for such highly symmetric matrices essentially reduces to verifying a finite set of linear inequalities. In fact, we show that in most cases they come from the same types of simple applications of Cauchy-Schwarz that are used in the Large sieve and the Larger Sieve. This allows us to define a simpler semidefinite relaxation which still retains some of the power of the full semidefinite relaxation introduced in this paper, and which is significantly more practical for numerical computations.

Definition 3. If $i, j \in Z$, then the expression $(j-i, \mathcal{P})$ is defined to be the set of partitions p in \mathcal{P} such that i and j lie in the same part of p .

Definition 4. For $d \subseteq \mathcal{P}$ we define $\phi(d)$ to be $\prod_{p \in d} (|p| - 1)$, and we set $\mu(d) = (-1)^{\#\{p \in d\}}$. For $s \subseteq \mathcal{P}$, we define $\phi_\kappa^s(d)$ by $\phi_\kappa^s(d) = \prod_{p \in d \cap s} (|p| - \kappa_p - 1) \prod_{p \in d \setminus s} (|p| - 1)$.

Theorem 3. *If a matrix B is given by*

$$B_{ij} = b_{(j-i, \mathcal{P})}$$

for a given system of real numbers b_d , then $B \succeq 0$ if the numbers b_d satisfy the system of inequalities

$$\forall k \subseteq \mathcal{P}, \sum_d \phi(\mathcal{P} \setminus (d \cup k)) \mu(k \setminus d) b_d \geq 0.$$

The above system of inequalities is satisfied iff there exist weights $w_k \geq 0$ such that

$$b_d = \sum_k \phi(d \cap k) \mu(k \setminus d) w_k,$$

or equivalently such that

$$B = \sum_k w_k \otimes_{p \in k} (|p|I_p - J_p) \otimes_{q \notin k} J_q.$$

If the natural map $Z \rightarrow \prod_{p \in \mathcal{P}} p$ is surjective, then $B \succeq 0$ if and only if the above conditions are satisfied.

Proof. We just need to prove this in the case $Z = \prod_{p \in \mathcal{P}} p$. We identify each partition p with a cyclic group $\mathbb{Z}/|p|$, and Z with the product group $\prod_{p \in \mathcal{P}} \mathbb{Z}/|p|$. We define the matrix $U \in \mathbb{R}^{Z \times \hat{Z}}$ to have as columns the set of characters $\chi : Z \rightarrow \mathbb{C}^\times$, so that $U_{i,\chi} = \chi(i)$. The matrix B is positive semidefinite if and only if the matrix $U^H B U \in \mathbb{R}^{\hat{Z} \times \hat{Z}}$ is positive semidefinite.

The ψ, χ entry of $U^H B U$ is given by

$$\begin{aligned} (U^H B U)_{\psi, \chi} &= \sum_{i,j} B_{ij} \overline{\psi(i)} \chi(j) = \sum_{i,j} b_{(j-i, \mathcal{P})} \overline{\psi(i)} \chi(j) \\ &= \sum_{x \in Z} b_{(x, \mathcal{P})} \sum_i \overline{\psi(i)} \chi(i+x) = \sum_{x \in Z} b_{(x, \mathcal{P})} \chi(x) |Z| \langle \psi, \chi \rangle. \end{aligned}$$

In particular, the matrix $U^H B U$ is diagonal, so in order to check that it is positive semidefinite we just need to check that the diagonal entries are nonnegative.

Every character χ on Z can be written as a product of characters χ_p on $\mathbb{Z}/|p|$. Letting k_χ be the set of p such that χ_p is nontrivial, we see that the χ, χ entry of $U^H B U$ is proportional to

$$\begin{aligned} \sum_{x \in Z} b_{(x, \mathcal{P})} \chi(x) &= \sum_{d \subseteq \mathcal{P}} b_d \sum_{(x, \mathcal{P})=d} \chi(x) = \sum_d b_d \prod_{p \in \mathcal{P} \setminus d} \begin{cases} |p| - 1 & p \notin k_\chi \\ -1 & p \in k_\chi \end{cases} \\ &= \sum_d b_d \phi(\mathcal{P} \setminus (d \cup k_\chi)) \mu((\mathcal{P} \setminus d) \cap k_\chi). \end{aligned}$$

Noting that the system of inequalities which we require the b_d s to satisfy has the same number of inequalities as variables, it is natural to look for a positive basis for the cone of solutions. By the multiplicative nature of these inequalities, it is enough to understand the case of just one partition p , in which case we see that a positive basis for the cone of solutions is given by $|p|I_p - J_p$ and J_p . Concretely, we can verify that the matrices

$$\bigotimes_{p \in k} (|p|I_p - J_p) \bigotimes_{q \notin k} J_q$$

form a positive basis for the cone of solutions by verifying that for any $k, k' \subseteq \mathcal{P}$, we have

$$\sum_d \phi(\mathcal{P} \setminus (d \cup k')) \mu(k' \setminus d) \cdot \phi(d \cap k) \mu(k \setminus d) = \begin{cases} 0 & k \neq k', \\ \prod_{p \in \mathcal{P}} (|p| - 2) & k = k'. \end{cases} \quad \square$$

Applying the above result to the case where some parts of some partitions have been knocked out, we have the following result.

Corollary 3. *If B^s is given by $B_{ij}^s = b_{(j-i, \mathcal{P})}^s$, then to check that $B^s \in \mathcal{B}_{s, \kappa}$ it is sufficient to check that the system of linear inequalities*

$$\forall k \subseteq \mathcal{P}, \sum_d \phi_\kappa^s(\mathcal{P} \setminus (d \cup k)) \mu(k \setminus d) b_d^s \geq 0$$

is satisfied. This occurs iff there are weights $w_k^s \geq 0$ such that

$$B^s = \sum_k w_k^s \bigotimes_{p \in k \cap s} ((|p| - \kappa_p) I_p - J_p) \bigotimes_{q \in k \setminus s} (|q| I_q - J_q) \bigotimes_{r \notin k} J_r,$$

and in this case the b_d^s s are given by

$$b_d^s = \sum_k \phi_\kappa^s(d \cap k) \mu(k \setminus d) w_k.$$

Examining the matrix

$$\bigotimes_{p \in k \cap s} \left((|p| - \kappa_p) I_p - J_p \right) \bigotimes_{q \in k \setminus s} \left(|q| I_q - J_q \right) \bigotimes_{r \notin k} J_r,$$

we see that there is little point in introducing a positive weight w_k^s for it if $k \setminus s$ is nonempty, since $|q| I_q - J_q$ is a positive combination of the matrices $(|q| - \kappa_q) I_q - J_q$ and J_q . This justifies considering the following simpler semidefinite relaxation.

Definition 5. If Z is a set, \mathcal{P} is a collection of partitions of Z , κ_p are integers attached to the partitions $p \in \mathcal{P}$, and D_s, D_f are collections of subsets of \mathcal{P} , then we define the D_s, D_f -relaxation of the sifting problem to be the problem of finding the following minimum λ such that there exist $B^d \in \mathcal{B}_{d, \kappa}$ for $d \in D_s$ and $w_d \geq 0$ for $d \in D_f$ with

$$\sum_{d \in D_s} B^d + \sum_{d \in D_f} w_d \bigotimes_{p \in d} \left((|p| - \kappa_p) I_p - J_p \right) \bigotimes_{q \notin d} J_q + J = \lambda I.$$

If $0 \in D_s$, then we also allow a matrix B^0 to be included, where the only constraint on B^0 is that its entries are nonnegative.

The Large Sieve and the Larger Sieve can both be thought of as living in this framework, with $D_s = \{1\}$ for the Large Sieve and $D_s = \{0\}$ for the Larger Sieve. In the case of the Larger Sieve, D_f contains only primes, and the corresponding weights w_p are proportional to $\log(p)$. One weakness of taking $D_s = \{0\}$ as in the Larger Sieve is that if the set Z contains any pair of elements $i \neq j$ such that $(i - j, \mathcal{P}) = \mathcal{P}$, then it will be impossible to find a system of weights $w_d \geq 0$ such that

$$\sum_{d \in D_f} w_d \bigotimes_{p \in d} \left((|p| - \kappa_p) I_p - J_p \right) \bigotimes_{q \notin d} J_q + J$$

has nonpositive off-diagonal entries, since the i, j entry will automatically be at least 1.

In the case of the Large Sieve, if only prime ds are used in D_f (as is common in many applications), then the weights w_d are proportional to $\frac{p - \kappa_p}{p}$.

The special case $D_s = \{1\}$ is particularly interesting in the case Z is an interval: if we have

$$B^1 + \sum_{d \in D_f} w_d \bigotimes_{p \in d} \left((|p| - \kappa_p) I_p - J_p \right) \bigotimes_{q \notin d} J_q + J = \lambda I,$$

then B^1 will have constant diagonal entries, so it will be a symmetric Toeplitz matrix, and can potentially be checked for positive semidefiniteness with a specialized algorithm. If λ is taken to be minimal, then B^1 will also be singular. In this case, we can apply the Pisarenko harmonic decomposition ([8], which uses Chapter 4 of [2]) to write

$$B^1 = V^H D V,$$

where D is a diagonal matrix with positive entries and dimension equal to the rank of B^1 , and each row of V has the form $(1, e^{2\pi i f}, \dots, e^{2\pi i (N-1)f})$ for some frequency f . The set of frequencies is determined by B^1 via the fact that B^1 and V must have the same nullspace. This decomposition of B^1 , together with the weights w_d , can be used to give very short proofs of upper bounds on the sizes of admissible tuples contained in short intervals.

7. THE FIRST LEVEL OF THE SEMIDEFINITE HIERARCHY

The Large Sieve is often stated in a weaker form than the one given here, where we restrict the set of d s considered to the primes, and this weaker form gives comparable results. Additionally, the Larger Sieve makes no use of non-prime d whatsoever, and in fact it is hard to modify it to take them into account in any useful way. Do we really need to use the matrices B^d with d non-prime to get good results? What happens if we restrict ourselves to just using the primes?

A preliminary result is that if p corresponds to a partition into just two parts, then there is never any point to consider $d = pk$ with k nontrivial. As a consequence, if all partitions have two parts, then the semidefinite framework considered here gets the exact answer at just the first level of the semidefinite hierarchy (i.e., using only B^p s with p corresponding to a single partition).

Theorem 4. *Suppose that p is a partition into two parts c_0, c_1 , and that D is a collection of subsets of $\mathcal{P} \setminus \{p\}$ such that the expression $(*)$, when restricted to c_i , proves the bound $|X \cap c_i| \leq \lambda_i$. Then using the collection of subsets $D \cup \{p\}$, the expression $(*)$ proves the bound $|X| \leq \max(\lambda_0, \lambda_1)$.*

Proof. Suppose that the bound on $|X \cap c_i|$ is proved via a system of matrices B_i^d . Then we can use the following matrices to prove the desired bound on $|X|$:

$$B^p = \begin{bmatrix} 0 & -J \\ -J & 0 \end{bmatrix}, \quad B^d = \begin{bmatrix} B_0^d & 0 \\ 0 & B_1^d \end{bmatrix},$$

where the indices are assumed to be arranged so that c_0 comes before c_1 , and the matrices $-J$ in B^p are rectangular rather than square. Then we have

$$B^p + \sum_{d \in D} B^d + J = \begin{bmatrix} \sum B_0^d + J & 0 \\ 0 & \sum B_1^d + J \end{bmatrix},$$

and the operator norm of this matrix is clearly equal to $\max(\lambda_0, \lambda_1)$. □

Now we consider a toy scenario where the partitions are “orthogonal”, so $Z = \prod_{p \in \mathcal{P}} p$. In this case it is easy to compute the true upper bound on $|X|$ (i.e., $|X| \leq \phi(\mathcal{P})$) using a symmetry argument, but the linear sieve framework is unable to prove this bound without using sieve weights λ_d supported on all subsets $d \subseteq \mathcal{P}$. In the semidefinite framework, the situation is nicer: we only need to use matrices B^p for p corresponding to partitions, and in fact the space of systems of matrices B^p that prove the true upper bound is enormous.

Theorem 5. *If $Z = \prod_{p \in \mathcal{P}} p$, then there is a high-dimensional set of systems $B^p \in \mathcal{B}_p$ such that*

$$\sum_{p \in \mathcal{P}} B^p + J = \phi(\mathcal{P})I.$$

Proof. Due to the symmetry of the problem, we only need to consider matrices B^p which have the property that B_{ij}^p only depends on $(i - j, \mathcal{P})$. So suppose that each B^p has the form

$$B_{ij}^p = b_{(i-j, \mathcal{P})}^p.$$

By Theorem 3, if we define $\phi^p(d)$ by

$$\phi^p(d) = \begin{cases} \phi(d) & p \notin d, \\ \frac{|p|-2}{|p|-1} \phi(d) & p \in d, \end{cases}$$

then in order for the system of matrices B^p to prove the bound $|X| \leq \phi(\mathcal{P})$, we need the system of numbers b_d^p to satisfy the following system of equations and inequalities:

$$\begin{aligned}
(1) \quad & \sum_d \phi^p(\mathcal{P} \setminus (d \cup k)) \mu(k \setminus d) b_d^p \geq 0 \text{ for all } p, k, \\
(2) \quad & \sum_p b_d^p + 1 = 0 \text{ for all } d \neq \mathcal{P}, \\
(3) \quad & \sum_p b_{\mathcal{P}}^p + 1 = \phi(\mathcal{P}).
\end{aligned}$$

Summing the first inequality above over $k' \subseteq k$ with weight $\phi^p(k')$ for some fixed k , we get

$$\sum_{d \supseteq k} \phi^p(\mathcal{P} \setminus d) b_d^p \geq 0 \text{ for all } p, k.$$

Summing the above inequality over $p \notin k$ with weights $\phi^{\mathcal{P}}(\mathcal{P} \setminus pk) \phi(\mathcal{P} \setminus p)^{-1}$ where $\phi^{\mathcal{P}}(k) = \prod_{q \in k} (|q| - 2)$, we get

$$\begin{aligned}
0 &\leq \sum_{p \notin k} \phi^{\mathcal{P}}(\mathcal{P} \setminus pk) \phi(\mathcal{P} \setminus p)^{-1} \sum_{d \supseteq k} \phi^p(\mathcal{P} \setminus d) b_d^p \\
&= \sum_{p \notin d} (b_{pd}^p + (|p| - 2)b_d^p) \phi(d)^{-1} \sum_{k \subseteq d} \phi^{\mathcal{P}}(\mathcal{P} \setminus pk) \\
&= \sum_d \phi^{\mathcal{P}}(\mathcal{P} \setminus d) \sum_p b_d^p \\
&= \left(\sum_p b_{\mathcal{P}}^p + 1 \right) - \sum_d \phi^{\mathcal{P}}(\mathcal{P} \setminus d) \\
&= \left(\sum_p b_{\mathcal{P}}^p + 1 \right) - \phi(\mathcal{P}) = 0,
\end{aligned}$$

so all the inequalities with $p \notin k$ must be satisfied with equality. This implies the system of equations

$$b_{pd}^p + (|p| - 2)b_d^p = 0 \text{ for all } p \notin d.$$

If we impose these extra equations, then the inequalities $\sum_{d \supseteq k} \phi^p(\mathcal{P} \setminus d) b_d^p \geq 0$ with $p \notin k$ will be satisfied with equality, and the equations $\sum_p b_d^p + 1 = 0$ for $d \neq \mathcal{P}$, $\sum_p b_{\mathcal{P}}^p + 1 = \phi(\mathcal{P})$ will become linearly dependent modulo these extra equations.

For the remaining inequalities, we introduce variables w_k^p such that

$$b_d^p = \sum_k \phi^p(d \cap k) \mu(k \setminus d) w_k^p,$$

and the remaining inequalities (1) will be satisfied iff we have $w_k^p \geq 0$ for all p, k . The system of identities $b_{pd}^p + (|p| - 2)b_d^p = 0$ for all $p \notin d$ is equivalent to the system of identities $w_k^p = 0$ for all $p \notin k$, since both are equivalent to (1) being satisfied with equality for all $p \notin k$.

Summing (2) and (3) with weights $\phi(\mathcal{P} \setminus (d \cup j))\mu(j \setminus d)$ and replacing b_d^p s by their definitions in terms of the w_k^p s, we get

$$\begin{aligned} \phi(\mathcal{P}) - \sum_d \phi(\mathcal{P} \setminus (d \cup j))\mu(j \setminus d) &= \sum_{d,p} \phi(\mathcal{P} \setminus (d \cup j))\mu(j \setminus d)b_d^p \\ &= \sum_{k,d,p} \phi(\mathcal{P} \setminus (d \cup j))\mu(j \setminus d)\phi^p(d \cap k)\mu(k \setminus d)w_k^p \\ &= \sum_{p,k} w_k^p \sum_d \phi(\mathcal{P} \setminus (d \cup j))\mu(j \setminus d)\phi^p(d \cap k)\mu(k \setminus d) \\ &= \sum_{p \in j} w_j^p \frac{|p| - 1}{|p|} |\mathcal{P}| - \sum_{p \notin j} w_{pj}^p \frac{1}{|p|} |\mathcal{P}|. \end{aligned}$$

Thus, for $j \neq \emptyset$ we have

$$\sum_{p \in j} w_j^p \frac{|p| - 1}{|p|} = \prod_p \frac{|p| - 1}{|p|} + \sum_{p \notin j} w_{pj}^p \frac{1}{|p|},$$

and for $j = \emptyset$ we have

$$1 = \prod_p \frac{|p| - 1}{|p|} + \sum_p w_p^p \frac{1}{|p|},$$

and this last equation follows from the equations for $j \neq \emptyset$ due to the linear dependence found earlier. We can now pick positive values for the w_j^p s by a downward induction on j , starting with $j = \mathcal{P}$, and at each step of the inductive construction we find that we have a high dimensional set of choices. \square

REFERENCES

- [1] Patrick X Gallagher. A larger sieve. *Acta Arithmetica*, 18(1):77–81, 1971.
- [2] Ulf Grenander, Gabor Szegő, and Mark Kac. Toeplitz forms and their applications. *Physics Today*, 11:38, 1958.
- [3] L. G. Hačijan. Polynomial algorithms in linear programming. *Zh. Vychisl. Mat. i Mat. Fiz.*, 20(1):51–68, 260, 1980.
- [4] Yu. V. Matiyasevich. *Desyataya problema Gilberta*, volume 26 of *Matematicheskaya Logika i Osnovaniya Matematiki [Monographs in Mathematical Logic and Foundations of Mathematics]*. VO “Nauka”, Moscow, 1993.
- [5] H. L. Montgomery. A note on the large sieve. *Journal of the London Mathematical Society*, s1-43(1):93–98, 1968.
- [6] Hugh L. Montgomery. The analytic principle of the large sieve. *Bull. Amer. Math. Soc.*, 84(4):547–567, 1978.
- [7] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- [8] V. F. Pisarenko. The Retrieval of Harmonics from a Covariance Function. *Geophysical Journal International*, 33(3):347–366, 09 1973.
- [9] Atle Selberg. *Collected papers. Vol. II*. Springer-Verlag, Berlin, 1991. With a foreword by K. Chandrasekharan.

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 77 MASSACHUSETTS AVENUE, BUILDING 2, ROOM 350B, CAMBRIDGE, MA 02139-7307
E-mail address: notzeb@mit.edu