CS 880: Advanced Complexity Theory 3/31/2008

Lecture 24: Additive Combinatorics

Instructor: Dieter van Melkebeek Scribe: Scott Diehl

Today we present an application of harmonic analysis in a non-Boolean setting. More specifically, we focus on additive combinatorics, which studies the structure of finite sets of integers or other abelian groups such as $\mathbb{Z}_m^n$ under addition.

1 Arithmetic Progressions

A typical question in additive combinatorics is as follows: for a subset A of the integers, how large is the set $A + A = \{a_1 + a_2 | a_1, a_2 \in A\}$? Trivially, we know that

$$|A| \leq |A + A| \leq |A|^2.$$

A random A yields $A + A$ of size closer to the high end of $|A|^2$. Furthermore, any set A whose members are sufficiently separated has $|A + A|$ close to this high end, and the high end is actually tight for sets A that are geometric progressions, e.g. the first n powers of two.

What about the low end of $|A|$: what is the structure of sets A where $|A+A|$ is close to $|A|$? We can see that this low-end estimate is essentially tight for sets A that are *arithmetic progressions*, e.g., the first n positive numbers.

This brings us to another question asked by additive combinatorics: do large (high-density) sets A have to contain long arithmetic progressions? For example, a recent result shows that the set of primes — which has nontrivial density by the prime number theorem — contains arithmetic progressions of arbitrary length.

Another result shows how large a set has to be to guarantee an arithmetic progression of length k. For an additive group G, let $r_k(G)$ be the maximum size of a subset $A \subseteq G$ such that A does not contain an arithmetic progression of length k. It was conjectured that any positive-density subset of the integers from 1 to N has an arithmetic progression of length k, i.e., that

$$r_k(\{1, \ldots, N\}) = o(N).$$

This conjecture has been resolved, and in fact, something stronger has been shown:

$$r_k(\{1, \ldots, N\}) \leq N/(\log \log N)^{c_k},$$

for some constant c_k depending only on k. In other words, density of only polynomial in $1/(\log \log N)$ suffices to guarantee an arithmetic progression of length k, where the degree of the polynomial depends on k.

The above statement is more general than the setting that we consider today. We focus on the special case of arithmetic progressions of length $k = 3$. The density known to guarantee such a progression is even smaller than in the general case, namely

$$r_3(\{1, \ldots, N\}) \leq cN\sqrt{\frac{\log \log N}{\log N}}, \tag{1}$$

where c is a constant.

Today, rather than on subsets of the integers, we focus on the additive group $\mathbb{Z}_p^n$ for primes $p > 2$, i.e., the set of vectors of length n over $\mathbb{Z}_p$. We prove a density result for this set similar to those cited above.

Theorem 1. *For any prime $p > 2$, there exists a constant $c_p > 0$ such that for all n,*

$$r_3(\mathbb{Z}_p^n) \leq c_p \frac{p^n}{n}.$$

In other words, any subset of $\mathbb{Z}_p^n$ of density $\Omega(1/n) = \Omega(1/\log N)$ has an arithmetic progression of length three, where $N = p^n = |\mathbb{Z}_p^n|$ and the hidden constant depends on p. We point out that the restriction that $p \neq 2$ is in place simply because there are no arithmetic progressions of length three over $\mathbb{Z}_2^n$, since the first element always ends up the same as the third. Also, the integer result of (1) follows by an involved argument from Theorem 1, although we will not cover this.

The proof of Theorem 1 involves harmonic analysis over $\mathbb{Z}_p^n$ for $p > 2$. Our previous results in this class involved harmonic analysis over $\mathbb{Z}_2^n$. For this result we will use the extension of these techniques to the non-Boolean setting, as discussed in Lecture 3.

2 Proof of Theorem 1

We use the probabilistic method to prove Theorem 1: For any fixed $A \subseteq \mathbb{Z}_p^n$ of sufficiently high density, we show that by picking three elements in a certain way, there is a positive probability that they form an arithmetic progression in A.

Fix $A \subseteq \mathbb{Z}_p^n$. We want to prove that if $\mu(A) > c_p/n$ then A contains an arithmetic progression of length 3, i.e., $(\exists x, y, z \in A)\, x \neq y,$ and $z = y + (y - x) = -x + 2y$.

Observe that if the measure of A is large enough, say, $2/3 + \epsilon$ for positive ϵ, then the theorem follows by a union bound. To see this, imagine picking $x, y \in \mathbb{Z}_p^n$ independently at random. We define success to be when all of x, y, and $z = -x + 2y$ fall inside A. Note that x, y, z form an arithmetic progression in A iff we have success and $x \neq y$. To calculate the probability of success, notice that each of x, y, and z are individually uniformly distributed, so the probability that each x, y, z individually lies outside of A is $\mu(\overline{A})$. By a union bound, the probability that not all of them lie inside A is at most $3\mu(\overline{A}) \leq 1 - 3\epsilon$. We lose at most another term of $1/p^n$ from the success probability to account for the case $x = y$. Therefore,

$$\Pr[x, y, z \text{ form an arithmetic progressio in } A] \geq 3\epsilon - 1/p^n.$$

This yields the desired positive probability for large enough n.

The simple union bound approach only works for densities larger than $2/3$, which is far from what we are aiming for. To do better, we find an exact expression for the success probability in terms of the Fourier expansion of the characteristic function of the set A. We focus on showing that the probability of success is larger than $1/p^n$.

To this end, define the function f to be

$$f(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise} \end{cases}$$

We claim that the success probability can be characterized in terms of the Fourier expansion of f as follows:

Claim 1.

$$\Pr[\textit{success}] = \sum_{s\in\Omega} (\hat{f}(s))^2 \hat{f}(-2s).$$

Proof. Recall that for $\Omega = \mathbb{Z}_p^n$, the characters of our Fourier expansion are $\chi_s(x) = \omega^{s\cdot x}$ for $s \in \Omega$, where $w = e^{2\pi i/p}$ and $s \cdot x$ denotes the vector inner product mod p. The characters χ_s are orthonormal with respect to $\langle f, g\rangle = \mathrm{E}[f(x)\overline{g(x)}]$ (notice the presence of the complex conjugate over g, which disappears in the Boolean case). Therefore, we can write

$$f = \sum_{s\in\Omega} \hat{f}(s)\chi_s, \text{ where } \hat{f}(s) = \langle f, \chi_s\rangle.$$

In particular, $\hat{f}(0) = \mathrm{E}[f] = \mu(A)$ for our particular choice of f as the characteristic sequence of A. Similarly, Parseval's equality allows us to relate the sum-of-squares of Fourier coefficients to $\mu(A)$:

$$\mu(A) = \langle f, f\rangle = \sum_s |\hat{f}(s)|^2.$$

Along the lines of some of our previous Fourier analysis, we write the success probability in terms of the Fourier expansion as follows:

$$\begin{aligned}\Pr[\text{success}] &= \mathrm{E}_{x,y}[f(x)f(y)f(z)] \text{ where } z = -x + 2y \\ &= \sum_{s,t,u\in\Omega} \hat{f}(s)\hat{f}(t)\hat{f}(u)\mathrm{E}_{x,y}[\chi_s(x)\chi_t(y)\chi_u(z)]. \qquad (2)\end{aligned}$$

We rewrite $\chi_u(z)$ as

$$\chi_u(z) = \chi_u(-x)\chi_u(2y) = \chi_{-u}(x)\chi_{2u}(y),$$

where the second inequality holds by definition of characters. This enables us to simplify the expectation:

$$\begin{aligned}\mathrm{E}_{x,y}[\chi_s(x)\chi_t(y)\chi_u(z)] &= \mathrm{E}_x[\chi_s\chi_{-u}] \cdot \mathrm{E}_y[\chi_t\chi_{2u}] \\ &= \langle\chi_s, \chi_u\rangle \cdot \langle\chi_t, \chi_{-2u}\rangle,\end{aligned}$$

where we switch χ_{-u} to χ_u and χ_{2u} to χ_{-2u} in the second equality in order to account for the complex conjugate in the definition of $\langle\cdot,\cdot\rangle$.

We now plug this expression into (2) and use the orthonormality of characters to conclude that:

$$\Pr[\text{success}] = \sum_{u\in\Omega} (\hat{f}(u))^2 \hat{f}(-2u).$$

□

So our goal is now to show that the expression in Claim 1 is greater than $1/|\Omega|$ for $\mu(A) \geq c_p/n$. Our strategy is to single out the null term. By our observations about $\hat{f}(0)$ above, we have

$$\Pr[\text{success}] = \sum_{s\in\Omega} (\hat{f}(s))^2 \hat{f}(-2s) = (\mu(A))^3 + \underbrace{\sum_{s\neq 0} (\hat{f}(s))^2 \hat{f}(-2s)}_{(*)}.$$

Therefore, we can bound our success probability by bounding $(*)$. Notice that we are done if none of the $\hat{f}(-2s)$ terms are large, since the sum-of-squares of coefficients is at most $\mu(A)$:

$$\sum_{s\in\Omega}(\hat{f}(s))^2\hat{f}(-2s) \geq (\mu(A))^3 - \max_{s\neq 0}(|\hat{f}(-2s)|)\underbrace{\sum_{u\neq 0}|\hat{f}(u)|^2}_{\leq\mu(A)}$$

$$\geq (\mu(A))^3 - \mu(A)\underbrace{\max_{s\neq 0}|\hat{f}(s)|}_{(**)}.$$

In order to bound this expression away from $1/|\Omega|$, we need to show that $(**)$ is small compared to $\mu(A)$. This presents us with the following dichotomy:

- Case 1: All $|\hat{f}(s)|$ for $s \neq 0$ are small compared to $\mu(A)$.
  Then $(**)$ is small with respect to $\mu(A)$ and we are essentially done.

- Case 2: For some $s \neq 0$, $|\hat{f}(s)|$ is large with respect to $\mu(A)$.
  This means that f correlates well with χ_s, i.e., χ_s is a good predictor of membership in A. As the different levels of χ_s define hyperplanes with perpendicular s, this means that there exists a hyperplane $H : x \cdot s = a$ for some $a \in \mathbb{Z}_p$ such that

  $$\Pr[x \in A | x \in H] \gg \mu(A).$$

  We claim that this reduces the problem to a smaller instance of the same problem, in particular, one on one fewer dimension. We know that H can be obtained as the bijective image of some affine transformation of the $(n-1)$-dimensional space, i.e., of some affine invertible $T : \mathbb{Z}_p^{n-1} \to H$. Then x, y, z is an arithmetic progression in H if and only if $T^{-1}(x), T^{-1}(y), T^{-1}(z)$ is an arithmetic progression in $\mathbb{Z}_p^{n-1}$. This follows immediately from linearity in the case where T is linear; the case of general affine transformations T follows since adding a constant to all elements of a sequence does not affect whether the sequence forms an arithmetic progression. Thus, in this case we can reduce the problem instance to one with larger density, $\mu' > \mu$, and one fewer variable, $n' = n - 1$.

In our analysis, we apply case 2 until the condition for case 2 fails to hold; we then apply case 1, which completes the proof. The density increases enough in each application to indicate that this process finishes before we run out of variables, since we have an upper bound of one on the density. Note, though, that we have to make sure that at each step our new density μ' satisfies the requirement that it is at least c_p/n', which becomes more stringent as n' decreases.

We now fill in the details, using a threshold of half of $\mu(A)^2$ for $(**)$ to be "large" or "small":

- Case 1: If $\max_{s\neq 0}|\hat{f}(s)| \leq \mu(A)^2/2$, then

  $$\Pr[success] \geq \mu^3/2.$$

- Case 2: $(\exists s \neq 0)\,|\hat{f}(s)| \geq \mu(A)^2/2$. To find a hyperplane with high correlation with membership in A, consider $g = f - \mu(A)$. We can write

  $$\mathrm{E}_{a\in\mathbb{Z}_p}[|\mathrm{E}_x[g(x) \mid s \cdot x = a]|] \geq |\hat{g}(s)| = |\hat{f}(s)| \geq \mu^2/2. \quad (3)$$

The first inequality holds because for uniform $x \in \mathbb{Z}_p^n$ the distribution of $s \cdot x$ is uniform over $\mathbb{Z}_p$, so

$$\begin{aligned} |\hat{g}(s)| = |\mathrm{E}_x[g(x)\overline{\chi_s(x)}]| &= |\mathrm{E}_{a\in\mathbb{Z}_p}[\mathrm{E}_x[g(x)\omega^{-a} \mid s \cdot x = a]]| \\ &\leq \mathrm{E}_{a\in\mathbb{Z}_p}\left[|\mathrm{E}_x[g(x)\omega^{-a} \mid s \cdot x = a]|\right] \\ &= \mathrm{E}_{a\in\mathbb{Z}_p}\left[|\mathrm{E}_x[g(x) \mid s \cdot x = a]|\right]. \end{aligned}$$

The equality $\hat{g}(s) = \hat{f}(s)$ for $s \neq 0$ holds because f and g only differ in a constant. The final inequality in (3) holds by assumption.

We also have that

$$\mathrm{E}_x[g] = \mathrm{E}_{a\in\mathbb{Z}_p}[\mathrm{E}_x[g(x) \mid s \cdot x = a]] = 0. \tag{4}$$

Adding (3) and (4) gives

$$\mathrm{E}_{a\in\mathbb{Z}_p}\left[\mathrm{E}_x[g(x) \mid s \cdot x = a] + |\mathrm{E}_x[g(x) \mid s \cdot x = a]|\right] \geq \mu^2/2.$$

This allows us to conclude that there exists an $a \in \mathbb{Z}_p$ such that

$$\mathrm{E}_x[g(x) \mid s \cdot x = a] + |\mathrm{E}_x[g(x) \mid s \cdot x = a]| \geq \mu^2/2,$$

and therefore,

$$\mathrm{E}_x[g(x) \mid s \cdot x = a] \geq \mu^2/4.$$

Our desired hyperplane H is defined by $s \cdot x = a$, and we have density $\mu' \geq \mu + \mu^2/4$ on H.

By applying case 2 no more than $4/\mu_0$ times, where $\mu_0 = \mu(A)$, we double the density to $\mu' \geq 2\mu_0$. We repeat this, using half as many applications of case 2 to double the density again. Repeating this process eventually halts since the density at any step is at most one. The total number of applications of case 2 is no more than

$$\sum_{i=0}^{\infty} 4/(2^i\mu_0) = 8/\mu_0.$$

After this many iterations, we end up in case 1 with $\mu' \geq \mu_0$ and $n' \geq n_0 - 8/\mu_0$. To establish the existence of an arithmetic progression of length 3 in A, it suffices to argue that $(\mu')^3/2 > 1/|\Omega'| = 1/p^{n'}$. This condition is met provided that

$$\mu_0^3/2 \geq 1/p^{n_0 - 8/\mu_0}. \tag{5}$$

For the right-hand side to make sense, we must have that μ_0 is at least $\Omega(1/n_0)$, say $\mu_0 \geq \frac{c}{n_0}$. For that setting (5) is met if

$$\left(\frac{c}{n_0}\right)^3 > \frac{2}{p^{(1-8/c)n_0}},$$

which is equivalent to

$$(1 - 8/c)n \log p > \log(\frac{2}{c^3}) + 3 \log n,$$

which holds for sufficiently large c. The latter condition defines the constant $c_p = c$ in the statement of the theorem.