Discrete analogues of the Kakeya set conjecture: The cases of Finite Fields and $\mathbb{Z}/N\mathbb{Z}$ for square-free N

John Green

January 25, 2021

These notes are made to accompany a presentation given during Jonathan Hickman's course "Incidence Geometry: Continuous and Discrete" in the second semester of the 2020-2021 academic year at the University of Edinburgh. The goal of these notes is to give a combined presentation of the proofs of the Finite Field Kakeya set conjecture following the paper of Dvir, Kopparty, Saraf and Sudan, and the proof of the Kakeya set conjecture in $\mathbb{Z}/N\mathbb{Z}$ for square-free N given in the paper of Dhar and Dvir. We also refer to the paper of Hickman and Wright for a more extensive discussion of the discrete analogues of the Kakeya set conjecture and related problems. See the bibliography for details.

Motivation

The Kakeya set conjecture and its connections to other important problems in harmonic analysis are well known, but let us briefly discuss it so as to understand how we should formulate our discrete analogues.

A Kakeya set in $\mathbb{R}^n$ is a set K containing a line in each direction. The Kakeya set conjecture asks if every such set has Hausdorff dimension n. This would be true if for each $s < n$, the s-dimensional Hausdorff measure of K was non-zero. It would also be sufficient to show that the n-dimensional Hausdorff measure was non-zero, however, as is well known there are examples of Kakeya sets for which the n-dimensional measure is 0.

The underlying Additive Combinatorial structure of this problem has been observed and exploited by Bourgain and others. This provides some motivation for why we might be interested in discrete analogues, for instance we might consider replacing $\mathbb{R}$ with a Finite Field, or if we aren't concerned with preserving the field structure of $\mathbb{R}$, we could instead consider $\mathbb{Z}/N\mathbb{Z}$.

We'll mostly be interested in $\mathbb{Z}/N\mathbb{Z}$. The notion of a Kakeya set in this setting is discussed in the Hickman and Wright paper; skipping a precise formulation for now, let's ask what the correct analogue of the conjecture should be.

It isn't helpful to talk about Hausdorff dimension for finite sets. We could instead ask for precise lower bounds on the size of such sets, or, in connection with the Euclidean case, we could think of these as being discrete approximations to a Euclidean Kakeya set and look for some kind of lower bound that holds as $N \to \infty$, that perhaps starts to resemble what we would see in the Euclidean case.

For the sake of motivation, we will assume in some vague, non-precise way that our discrete Kakeya $K = K_N \subseteq (\mathbb{Z}/N\mathbb{Z})^n$ sets are approximating a Euclidean Kakeya set $K_{\mathbb{E}}$ (intersected with $[0,1]^n$, say) at appropriate scales. Suppose that $(\mathbb{Z}/N\mathbb{Z})^n$ represents a lattice of points separated by $1/N$ that approximates a Euclidean Kakeya set on scale $1/N$. Comparison with the s-dimensional Hausdorff measure approximated at scale $1/N$ shows that the natural approximating quantity is

$|K|/N^s$, and we are interested in lower bounds on this that are independent of N - a statement such as "$|K|/N^s \geq C$" can be interpreted as saying "The s-dimensional Hausdorff measure of K is at least C".
Continuing this analogy, we expect that it should be impossible to find estimates of the form $|K| \geq CN^n$ for some $C > 0$ independent of N, and in the paper of Hickman and Wright, it is shown that this is true. However, this is not the case in the Finite Field setting, as we shall see. We shall note why later - but for now, let us agree that the correct conjecture should be:

Conjecture. *For each $\varepsilon > 0$, there exists $C = C(n, \varepsilon)$ independent of N so that any Kakeya set K in $(\mathbb{Z}/N\mathbb{Z})^n$ satisfies $|K| \geq CN^{n-\varepsilon}$.*

In these notes, we will first give a proof of the sharp Finite Field Kakeya set bounds, and then extend this result to give a proof of the Kakeya set conjecture for square-free N, although we remark at this stage that the formulation of Kakeya set will be slightly different to that given in the Hickman and Wright paper in the latter case.

1 The Kakeya set conjecture in Finite Fields

We will first discuss the results in the Finite Field case due to Dvir, Kopparty, Saraf and Sudan. We will only need this result for the fields $\mathbb{Z}/p\mathbb{Z}$ for prime p, but the proof is the same in any finite field. Let $\mathbb{F}_q$ denote the field of size q. A Kakeya set in $\mathbb{F}_q^n$ is a set K for which given any non-zero vector $b \in \mathbb{F}_q^n$, there is an $a \in \mathbb{F}_q^n$ such that the line $L = \{a + tb : t \in \mathbb{F}_q\}$ is contained in K. For later reference, notice that we can identify direction vectors b which are the same up to scaling, and the set of these equivalence classes forms the projective space $\mathbb{PF}_q^{n-1}$.
We have the following lower bound:

Theorem 1. *If $K \subseteq \mathbb{F}_q^n$ is a Kakeya set, then $|K| \geq \frac{q^n}{(2-\frac{1}{q})^n}$.*

Immediately, we have that $|K| \geq C_n|\mathbb{F}_q|^n$ for any finite field.
Before we go into the proof of this theorem, we will need to review some facts about polynomials which will form the basis of the polynomial methods in this work.

1.1 Hasse derivatives and the extended Schwartz-Zippel lemma

In this section we introduce the Hasse derivative, a formal notion of derivative for polynomials replacing the classical derivative in polynomial rings over general fields. We then introduce a notion of multiplicity analogous to the usual notion of a polynomial vanishing to a given order, and prove a strengthened version of the Schwartz-Zippel lemma which will form the basis of the proof of Theorem 1.
First, let us introduce some notation. For multiindices $\alpha \in \mathbb{N}_0^n$ we will write $|\alpha| = \alpha_1 + \ldots + \alpha_n$. We will use upper case X and Y to denote vectors $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_m)$ where the lengths will be obvious from the context or stated otherwise. We will use $\mathbb{F}$ to denote a general field, and $\mathbb{F}[X]$ to denote the ring of polynomials over X in variables $x_1, \ldots, x_n$. We use X^α to denote $x_1^{\alpha_1} \ldots x_n^{\alpha_n}$. Given a polynomial $P(X)$ of degree d, we let $H_P(X)$ denote the sum of the terms of degree d, that is, H_P is the homogeneous part of P having highest total degree.
For multiindices α and β we denote

$$\binom{\alpha}{\beta} = \prod_{i=1}^{n} \binom{\alpha_i}{\beta_i}.$$

Note that this is the coefficient of $X^\beta Y^{\alpha-\beta}$ in the expansion of $(X+Y)^\alpha$.

Definition. Given $P \in \mathbb{F}[X]$, the α^{th} Hasse derivative of P, denoted $P^{(\alpha)}$, is the polynomial which is the coefficient of Y^α in the expansion of $P(X+Y)$, that is,

$$P(X+Y) = \sum_\alpha P^{(\alpha)}(X)Y^\alpha.$$

The multiplicity of P at a point A, denoted $\text{mult}(P,A)$, is defined to be the largest integer M for which $P^{(\alpha)}(A) = 0$ for all α with $|\alpha| < M$ (we set $\text{mult}(P,A) = \infty$ if there is no such largest integer). For vectors of polynomials $P = (P_1, \ldots, P_m) \in \mathbb{F}[X]^m$, set $\text{mult}(P,A) = \min_i\{\text{mult}(P_i,A)\}$.

Here are some basic properties of Hasse derivatives and multiplicities:

Proposition 1. *Let $P, Q \in \mathbb{F}[X]$, $\alpha, \beta \in \mathbb{N}_0^n$, $\lambda, \mu \in \mathbb{F}$. Then:*

1. *$P(A) = 0$ if and only if mult$(P,A) \geq 1$.*
2. *$\lambda P^{(\alpha)} + \mu Q^{(\alpha)} = (\lambda P + \mu Q)^{(\alpha)}$.*
3. *If P is homogeneous of degree d, then $P^{(\alpha)}$ is homogeneous of degree $d - |\alpha|$ or $P^{(\alpha)} = 0$.*
4. *$(H_P)^{(\alpha)} = H_{P^{(\alpha)}}$ or $(H_P)^{(\alpha)} = 0$.*
5. *$\left(P^{(\alpha)}\right)^{(\beta)} = \binom{\alpha}{\beta} P^{(\alpha+\beta)}$.*
6. *If $A \in \mathbb{F}^n$ is such that mult$(P,A) = m$, then mult$(P^{(\alpha)}, A) \geq m - |\alpha|$.*

Proof. We have that $\text{mult}(P,A) \geq 1$ if and only if $P^{(\alpha)}(A) = 0$ for each $|\alpha| < 1$ - that is, $\alpha = 0$. This proves the first statement.
The second statement follows from expanding $P(X+Y)$ and $Q(X+Y)$ and grouping coefficients, which is easily seen. The third statement follows from expanding $\lambda^d P(X+Y) = P(\lambda X + \lambda Y)$ in two different ways and comparing coefficients.
By the linearity just observed, we can write $P = H_P + R$, where the degree of R is strictly less than the degree of P, and we have $H_P^{(\alpha)} = P^{(\alpha)} - R^{(\alpha)}$. If this is non-zero, it must be homogeneous of degree $\deg P - |\alpha|$, hence $P^{(\alpha)} - R^{(\alpha)} = H_{P^{(\alpha)} - R^{(\alpha)}}$. However, the degree of $R^{(\alpha)}$ is strictly less than $\deg P - |\alpha|$, so we must have

$$P^{(\alpha)} = P^{(\alpha)} - R^{(\alpha)} = H_{P^{(\alpha)} - R^{(\alpha)}} = H_{P^{(\alpha)}}.$$

For the fifth statement, we expand $P(X+Y+Z)$ in two different ways. Firstly,

$$\begin{aligned} P(X+(Y+Z)) &= \sum_\alpha P^{(\alpha)}(X)(Y+Z)^\alpha \\ &= \sum_\alpha \sum_{\beta+\gamma=\alpha} P^{(\alpha)}(X) \binom{\alpha}{\beta} Y^\gamma Z^\beta \\ &= \sum_{\beta,\gamma} P^{(\beta+\gamma)}(X) \binom{\beta+\gamma}{\beta} Y^\gamma Z^\beta. \end{aligned}$$

Also, we may write

$$P((X+Y)+Z) = \sum_\beta P^{(\beta)}(X+Y)Z^\beta = \sum_\beta \sum_\gamma \left(P^{(\beta)}\right)^{(\gamma)}(X) Y^\gamma Z^\beta.$$

Comparing coefficients yields the result.
We now prove the final statement. By assumption, for any β with $|\beta| < m$, we have $P^{(\beta)}(A) = 0$. For each γ such that $|\gamma| < m - |\alpha|$, we have

$$(P^{(\alpha)})^{(\gamma)}(A) = \binom{\alpha+\gamma}{\gamma} P^{(\alpha+\gamma)}(A).$$

Since $|\alpha + \gamma| < m$, we have $(P^{(\alpha)})^{(\gamma)}(A) = 0$. Thus $\text{mult}(P^{(\alpha)}, A) \geq m - |\alpha|$. □

We will also need some basic results on the behaviour of multiplicities under composition of polynomials. Given $P \in \mathbb{F}[X]^m, Q \in \mathbb{F}[Y]^n$, let us consider the polynomial $P(Q(Y))$. We have the following:

Proposition 2. *For any A, $mult(P \circ Q, A) \geq mult(P, Q(A))mult(Q - Q(A), A)$. In particular, since $mult(Q - Q(A), A) \geq 1$,we have $mult(P \circ Q, A) \geq mult(P, Q(A))$.*

Proof. Let $m_1 = \text{mult}(P, Q(A))$ and $m_2 = \text{mult}(Q - Q(A), A)$. Note that $m_2 \geq 1$. If $m_1 = 0$ we are done, so assume $m_1 \geq 1$, so that $P(Q(A)) = 0$. Now we have

$$\begin{aligned}
P(Q(A+Y)) &= P\left(Q(A) + \sum_{\alpha \neq 0} Q^{(\alpha)}(A) Y^\alpha\right) \\
&= P\left(Q(A) + \sum_{|\alpha| \geq m_2} Q^{(\alpha)}(A) Y^\alpha\right) && \text{since } \text{mult}(Q - Q(A), A) = m_2 > 0 \\
&= P\left(Q(A) + R(Y)\right) && \text{where } R(Y) = \textstyle\sum_{|\alpha| \geq m_2} Q^{(\alpha)}(A) Y^\alpha \\
&= P(Q(A)) + \sum_{\beta \neq 0} P^{(\beta)}(Q(A)) R(Y)^\beta \\
&= \sum_{|\beta| \geq m_1} P^{(\beta)}(Q(A)) R(Y)^\beta && \text{since } \text{mult}(P, Q(A)) = m_1 > 0.
\end{aligned}$$

Since each monomial Y^α appearing in R has $|\alpha| \geq m_2$, and $R(Y)$ is raised to the power β with $\beta \geq m_1$, we conclude that $P(Q(A+Y))$ is of the form $\sum_{|\gamma| \geq m_1 m_2} c_\gamma Y^\gamma$, and the result follows. □

In practice we will only use this result to say that for $A, B \in \mathbb{F}^n$, the single variable polynomial $P_{A,B}(T) := P(A + TB)$ has $\text{mult}(P_{A,B}, t) \geq \text{mult}(P, A + tB)$ for each $t \in \mathbb{F}$.
We are now ready to prove the strengthened Schwartz-Zippel lemma.

Lemma 3. *Let $P \in \mathbb{F}[X]$ be a non-zero polynomial of degree at most d. Then for any finite $S \subseteq \mathbb{F}$, we have*

$$\sum_{A \in S^n} mult(P, A) \leq d|S|^{n-1}.$$

Proof. We induct on n. For $n = 1$, we must show that for a polynomial of one variable, the sum of multiplicities at each point of S is at most the degree d. Clearly, it is enough to show that if $\text{mult}(P, A) = m$ then $(X - A)^m$ divides P, so that P factors as $(\prod_{A \in S}(X - A)^{\text{mult}(P,A)})Q(X)$. In this case, we have that $P(A + Y) = \sum_\alpha P^{(\alpha)}(A) Y^\alpha$ and $P^{(\alpha)}(A) = 0$ for all $\alpha < m$. Thus Y^m divides $P(A + Y)$, and setting $Y = X - A$ concludes this case.

Now suppose $n > 1$. Write

$$P(x_1, \ldots, x_n) = \sum_{j=0}^{t} P_j(x_1, \ldots, x_{n-1}) x_n^j,$$

where $0 \leq t \leq d$, P_t is non-zero and $\deg P_j \leq d - j$. For any $a_1, \ldots, a_{n-1} \in S$, denote $m_{a_1,\ldots,a_{n-1}} = \text{mult}(P_t, (a_1, \ldots, a_{n-1}))$. We first show that

$$\sum_{a_n \in S} \text{mult}(P, (a_1, \ldots, a_n)) \leq m_{a_1,\ldots,a_{n-1}} |S| + t.$$

Let $\alpha \in \mathbb{N}_0^{n-1}$ be such that $|\alpha| = m_{a_1,\ldots,a_{n-1}}$ and $P_t^{(\alpha)} \neq 0$. Then we have that

$$P^{(\alpha,0)}(x_1, \ldots, x_n) = \sum_{j=0}^{t} P_j^{(\alpha)}(x_1, \ldots, x_{n-1}) x_n^j$$

and hence $P^{(\alpha,0)}$ is non-zero (since $P_t^{(\alpha)} \neq 0$). So by the previous proposition,

$$\begin{aligned} \text{mult}(P, (a_1, \ldots, a_n)) &\leq |(\alpha, 0)| + \text{mult}(P^{(\alpha,0)}(x_1, \ldots, x_n), (a_1, \ldots, a_n)) \\ &\leq m_{a_1,\ldots,a_{n-1}} + \text{mult}(P^{(\alpha,0)}(a_1, \ldots, a_{n-1}, x_n), a_n). \end{aligned}$$

Summing over $a_n \in S$, and applying the $n = 1$ case to $P^{(\alpha,0)}(a_1, \ldots, a_{n-1}, x_n)$ (which has degree t), we get the desired inequality.
We may now bound

$$\sum_{a_1,\ldots,a_n \in S} \text{mult}(P, (a_1, \ldots, a_n)) \leq \left(\sum_{a_1,\ldots,a_{n-1} \in S} m_{a_1,\ldots,a_{n-1}} \right) |S| + |S|^{n-1} t.$$

By the inductive hypothesis, the sum in brackets is bounded by $(d-t)|S|^{n-2}$, which completes the proof. □

This gives an important corollary in Finite Fields.

Corollary 1. *Let $P \in \mathbb{F}_q[X]$ be a polynomial of degree at most d. If $\sum_{A \in \mathbb{F}_q^n} mult(P, A) > dq^{n-1}$, then $P = 0$.*

1.2 Proof of the lower bound

In this section we prove Theorem 1. At this point we introduce some more notation to simplify some expressions that will appear in applying our polynomial methods. We use

$$\delta_{n,d} = \binom{d+n-1}{n-1} = \binom{d+n-1}{d}$$

to denote the dimension of the space of homogeneous polynomials of degree d in n variables over a given field, and

$$\Delta_{n,d} = \binom{d+n}{n} = \binom{d+n}{d}$$

to denote the dimension of the space of polynomials of degree at most d in n variables (these equalities follow from simple combinatorial arguments, e.g. stars and bars).
The following proposition expresses our basic polynomial technique:

Proposition 4. *Given a set $K \subseteq \mathbb{F}^n$ and non-negative integers m, d such that $\Delta_{n,m-1}|K| < \Delta_{n,d}$, there exists a non-zero polynomial $P \in \mathbb{F}[X]$ of total degree at most d such that mult$(P, A) \geq m$ for every $A \in K$.*

Proof. For a given A, the condition that mult$(P, A) \geq m$ means that the rank 1 linear function $P \mapsto P^{(\alpha)}(A) = 0$ for each multiindex α with $|\alpha| < m$. This imposes $\Delta_{n,m-1}$ linear constraints on P. Since the total number of linear constraints is $\Delta_{n,m-1}|K|$, which is strictly less than the dimension of the space of polynomials of degree at most d in n variables, so there is a non-zero polynomial of degree d vanishing to multiplicity m at every point of K. □

Proof of Theorem 1. Let l be a large multiple of q and let $m = 2l - l/q$, $d = lq - 1$. Note that $d < lq$ and thus $(m-l)q = ql - l > d - l$. We will prove by contradiction that $|K| \geq \Delta_{n,d}/\Delta_{n,m-1}$, so let us assume that $\Delta_{n,m-1}|K| < \Delta_{n,d}$. By the previous proposition, there exists a non-zero polynomial $P \in \mathbb{F}[X]$ of degree $d^* \leq d$ such that mult$(P, A) \geq m$ for each $A \in K$.
Note that $d^* \geq l$ since $m \geq l$ and, since P vanishes to multiplicity m at some A but is non-zero, there must be monomials of degree greater than m in the expansion of $P(A + (X - A))$.
We will show that H_P vanishes to multiplicity l at each point $B \in \mathbb{F}_q^n$. Let α be such that $|\alpha| < l$. Denote $Q = P^{(\alpha)}$, and let $d' \leq d^* - |\alpha|$ be the degree of Q. Pick A such that $\{A + tB : t \in \mathbb{F}_q\} \subseteq K$. Then for all $t \in \mathbb{F}_q$, mult$(Q, A + tB) \geq m - |\alpha|$.
Since $|\alpha| < l$ and $(m-l)q > d-l \geq d^*-l$, we get $(m-|\alpha|)q = (m-l)q+(|\alpha|-l)q > d^*-l+(|\alpha|-l)q = d^* - |\alpha| + (l - |\alpha|)(q - 1) > d^* - |\alpha|$.
Let $Q_{A,B}(T)$ be the polynomial $Q(A + TB)$. Then mult$(Q_{A,B}, t) \geq$ mult$(Q, A + tB) \geq m - |\alpha|$. Since $(m - |\alpha|)q > d^* - |\alpha| \geq d' \geq \deg Q_{A,B}$, the $n = 1$ case of the corollary of the Schwartz-Zippel lemma implies $Q_{A,B} = 0$.
Therefore the coefficient of $T^{d'}$ in $Q_{A,B}$ is 0. It is easily checked that this coefficient is equal to $H_Q(B)$, so $H_Q(B) = 0$. Thus $(H_P)^{(\alpha)}(B) = (H_Q(B)$ or $0) = 0$. Since this is true for all α with $|\alpha| < l$, we have mult$(H_P, B) \geq l$.
Note. For future reference, observe that we have shown for a homogeneous polynomial of degree at most $lq - 1$ in n variables that if it vanishes to multiplicity at least m at each point of a line L in direction B, then it vanishes to multiplicity at least l at B.
By the corollary of the Schwartz-Zippel lemma, noting that $lq^n > d^*q^{n-1}$, we conclude that $H_P = 0$, which in turn means $P = 0$, a contradiction. At this point all that remains is to tidy up our lower bound, using the freedom we have in our choice of l.
We have

$$|K| \geq \binom{d+n}{n} / \binom{m+n-1}{n} = \binom{lq-1+n}{n} / \binom{2l-l/q+n-1}{n}$$
$$= \frac{\prod_{i=1}^{n}(lq-1+i)}{\prod_{i=1}^{n}(2l-l/q-1+i)} = \frac{\prod_{i=1}^{n}(q-1/l+i/l)}{\prod_{i=1}^{n}(2-1/q-1/l+i/l)}$$

Since l can be any large multiple of q, we can take the limit $l \to \infty$ to get the result. □

2 Proof of the Kakeya set conjecture in $Z/N\mathbb{Z}$ for square-free N

In this case, we need to take a bit more care than the case of fields about how we define our Kakeya sets. But before we do that, let us briefly comment on why the Finite Field case has some inadequacies as being a model for the Euclidean case, and why $\mathbb{Z}/N\mathbb{Z}$ is a step in the right direction. This connection will provide some explanation as for why the Finite Field bound is better than

expected, and we will see how the improvements that $\mathbb{Z}/N\mathbb{Z}$ possesses relate to the "worse" lower bound in the more general Theorem.
The issue is that of the scales available in each case. In the Euclidean case, the usual distance gives a range of infinitely many scales which are ubiquitous in many arguments. In the Finite Field case, there are no natural notions of distance that provide any more scales than the trivial discrete distance. In the setting of rings such as $\mathbb{Z}/N\mathbb{Z}$, this is no longer the case - the divisors of N provide a range of scales to work with. For instance, we could consider the balls B_d in $(\mathbb{Z}/N\mathbb{Z})^n$, indexed by divisors of N, defined as $B_d := \{x = (x_1, \ldots, x_n) \in (\mathbb{Z}/N\mathbb{Z})^n : \|x\| \preceq d\}$ where $\|x\| := N/\gcd(x_1, \ldots, x_n, N)$ and $a \preceq b$ if and only if $a|b$. We will not develop this theory any further, but note that the divisors of N give a natural (finite) range of scales, and it is through this connection that we might hope to extend our arguments to the Euclidean case.
This should provide some motivation for why the bound expressed below is "natural" - in adding in more scales, that is, more divisors for N, we are obtaining a better model for the Euclidean case, and in the bound below, we do not obtain a lower bound of the form $C_n N^n$ unless we bound the number of divisors of N.

Theorem 2. *Let $N = p_1 \ldots p_r$ be a square-free integer with distinct prime factors $p_1, \ldots, p_r$. Then for each Kakeya set $K \subseteq (\mathbb{Z}/N\mathbb{Z})^n$, we have*

$$|K| \geq \frac{N^n}{\prod_{i=1}^r (2 - 1/p_i)^n}.$$

At this stage, we have not even been precise about what a Kakeya set in $(\mathbb{Z}/N\mathbb{Z})^n$ is. We will now go into these preliminary definitions and results.

2.1 Preliminaries

Recall the Chinese remainder theorem - if m and n are coprime, then $\mathbb{Z}/mn\mathbb{Z}$ and $\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$ are isomorphic as rings. Applied iteratively, one can say that for a square-free integer N having distinct prime factors $p_1, \ldots, p_r$, that $\mathbb{Z}/N\mathbb{Z}$ is isomorphic as a ring to $\mathbb{F}_{p_1} \times \ldots \times \mathbb{F}_{p_r}$, where $\mathbb{F}_{p_i}$ denotes the Finite Field $\mathbb{Z}/p_i\mathbb{Z}$. Thus we may identify $(\mathbb{Z}/N\mathbb{Z})^n$ with $\mathbb{F}_{p_1}^n \times \ldots \times \mathbb{F}_{p_r}^n$.
In the case of one finite field, we could take the definition of our set of directions to be the projective space $\mathbb{P}\mathbb{F}_{p_i}^{n-1}$. In general, denoting $R = \mathbb{Z}/N\mathbb{Z}$, we will take as our set of directions the projective space $\mathbb{P}R^{n-1} := \mathbb{P}\mathbb{F}_{p_1}^{n-1} \times \ldots \times \mathbb{P}\mathbb{F}_{p_r}^{n-1}$. It is a basic fact following from the Chinese remainder theorem that the set $\{tb : t \in \mathbb{Z}\}$, where $b \in \mathbb{P}R^{n-1}$, is independent of the choices of representative for each element of the $\mathbb{P}\mathbb{F}_{p_i}^{n-1}$.
Note that this definition of the projective space differs from that given in the Hickman and Wright paper, where it is instead asked that at least one entry in an n-tuple from $(\mathbb{Z}/N\mathbb{Z})^n$ be invertible. Through the Chinese remainder theorem, this would mean an element of $\mathbb{F}_{p_1}^n \times \ldots \times \mathbb{F}_{p_r}^n$ in which for some j, the j^{th} coordinate in each $\mathbb{F}_{p_i}^n$ is non-zero. In our definition, the "j" for which we have a non-zero entry in $\mathbb{F}_{p_i}^n$ can be different for each i, so our formulation contains more directions. This means that our Kakeya set should be "bigger" than those given in the Hickman and Wright formulation, however, this gives a more natural formulation for square-free N, as it does not depend on a choice of generating set for $\mathbb{Z}/N\mathbb{Z}$.
The proof of Theorem 2, roughly speaking, proceeds as follows. We can compare the size of Kakeya set to the rank of some matrix, which we will construct. In the case of $\mathbb{F}_p$, we could take this to be a "line matrix" with rows being the indicator vectors of some lines, one chosen for each direction, columns indexed by elements of $\mathbb{F}_p^n$. We then multiply this matrix by some matrix to get a matrix of possibly lower rank, but that is well understood and its rank easily bounded below. In the $\mathbb{F}_p$

case, this pairs with known results on the "point-hyperplane incidence matrix" to give a lower bound on the size of Kakeya sets in $\mathbb{F}_p$. In general, we will need to construct a more complicated matrix that combines the line matrix with a matrix constructed to multiply to the "right" matrix to get a lower bound on the rank, but once we unpack all the arguments, we see that this is only possible because of the same polynomial method techniques from the finite field case.
We will briefly need to make use of the line matrix in the general case, so let us discuss it here. Given a Kakeya set $S \subseteq R^n$, where $R = \mathbb{Z}/N\mathbb{Z}$, for each direction $b \in \mathbb{P}R^{n-1}$ choose a line $L(b) \subseteq S$ in direction b. Define the line matrix M_S with rows indexed by directions b and columns indexed by elements of R^n by setting its rows to be the indicator vectors of $L(b)$, that is, the entry corresponding to a direction b and point X is 1 if X lies on $L(b)$, 0 otherwise.

Proposition 5. *For any field $\mathbb{F}$, the line matrix M_S of a Kakeya set $S \in R^n$ has rank at least $|S'|/|R|$, where S' is the set of indices corresponding to non-zero columns of M_S. Also, S' is itself a Kakeya set.*

Proof. First pick a non-zero line $L_1 = L(b_1)$. Given lines $L_1 = L(b_1), \ldots, L_{t-1} = L(b_{t-1})$, the cardinality of their union is at most $|R|(t-1)$. If $|R|(t-1) < |S'|$, there is a column corresponding to a point which does not intersect $L_1, \ldots, L_{t-1}$, but does intersect some $L_t(b_t)$. Hence if we cannot add another such line to the collection, the final number t of lines satisfies $|R|t \geq |S'|$. Furthermore, by construction the row corresponding to b_t is independent over any field to those corresponding to $b_1, \ldots, b_{t-1}$, so the rank of M_S is at least t, and we obtain the result.
That S' is a Kakeya set is straightforward - given a direction b, consider the line $L(b)$ used to construct M_S. By construction, S' contains each point of this line, for whenever $L(b)$ hits a point, we have a 1 entry, so that column is non-zero. □

We will need some basic tools from Linear Algebra, which we shall now review. Firstly, we need the following dimension bound for tensors:

Proposition 6. *Let U and V be finite dimensional vector spaces over a field $\mathbb{F}$, let $u_1, \ldots, u_n \in U$ linearly independent, and for each $i \in \{1, \ldots, n\}$, let $v_1^{(i)}, \ldots, v_m^{(i)} \in V$ be linearly independent. Then the elementary tensors $u_i \otimes v_j^{(i)}$ form a linearly independent collection of size nm in $U \otimes V$.*

Proof. Let $w_1, \ldots, w_l$ be a basis of V and write $v_j^{(i)} = \sum_{k=1}^{l} \lambda_{i,j,k} w_k$. Note that the $u_i \otimes w_k$ form a linearly independent collection in $U \otimes V$. We must show that

$$\sum_{i=1}^{n} \sum_{j=1}^{m} \alpha_{i,j} u_i \otimes v_j^{(i)} = 0 \Rightarrow \forall i, j,\ \alpha_{i,j} = 0.$$

We have that

$$0 = \sum_{i=1}^{n} \sum_{j=1}^{m} \alpha_{i,j} u_i \otimes v_j^{(i)} = \sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{l} \alpha_{i,j} \lambda_{i,j,k} u_i \otimes w_k$$

and so by linear independence of the $u_i \otimes w_k$, we have that for each i, k, $\sum_{j=1}^{m} \alpha_{i,j} \lambda_{i,j,k} = 0$ and hence $\sum_{k=1}^{l} \sum_{j=1}^{m} \alpha_{i,j} \lambda_{i,j,k} w_k = \sum_{j=1}^{m} \alpha_{i,j} v_j^{(i)} = 0$ for each i. Thus by linear independence of the $v_j^{(i)}$, we have $\alpha_{i,j} = 0$. □

We shall also make use of the Kronecker product of two matrices.

Definition. Given an $m \times n$ matrix A and a $r \times s$ matrix B, the Kronecker product $A \otimes B$ is the $mr \times ns$ block matrix given by

$$\begin{bmatrix} a_{1,1}B & \cdots & a_{1,n}B \\ \vdots & \ddots & \vdots \\ a_{m,1}B & \cdots & a_{m,n}B \end{bmatrix}.$$

This product satisfies a number of helpful properties, however, here we shall only need the following:

Proposition 7. *Let $A = (a_{i,j})$ be an $m \times n$ matrix, $B = (b_{i,j})$ an $r \times s$ matrix, $X = (x_{i,j})$ an $n \times p$ matrix and $Y = (y_{i,j})$ an $s \times t$ matrix. Then $(A \otimes B)(X \otimes Y) = (AX) \otimes (BY)$.*

Proof. Let us index rows in $(A \otimes B)$ by pairs (i_1, i_2) corresponding to rows of A and B, and columns by pairs (j_1, j_2) corresponding to those of A and B, and similarly for the other Kronecker products. For instance, the $((i_1, i_2), (j_1, j_2))$ entry of $A \otimes B$ is $a_{i_1,j_1} b_{i_2,j_2}$. We have that the $((i_1, i_2), (k_1, k_2))$ entry of $(A \otimes B)(X \otimes Y)$ is given by

$$\sum_{(j_1,j_2)} a_{i_1,j_1} b_{i_2,j_2} x_{j_1,k_1} y_{j_2,k_2} = \left(\sum_{j_1} a_{i_1,j_1} x_{j_1,k_1} \right) \left(\sum_{j_2} b_{i_2,j_2} y_{j_2,k_2} \right).$$

The right hand side is precisely the $((i_1, i_2), (k_1, k_2))$ entry of $(AX) \otimes (BY)$. □

Finally, we will also need the concept of crank for a set of matrices.

Definition. Given a finite set $T = \{A_1, \ldots, A_n\}$ of matrices having the same number of columns we let $\mathsf{crank}(T)$ be the rank of the matrix obtained by concatenating all the elements A_i in T along their columns. Equivalently, it is the dimension of the space spanned by $\cup_{i=1}^{n} \{r : r \text{ is a row in } A_i\}$.

Some basic properties follow.

Proposition 8. *Given $m \times n$ matrices $A_1, \ldots, A_r$ and an $n \times p$ matrix B we have $\mathsf{crank}\{A_i\}_{i=1}^{r} \geq \mathsf{crank}\{A_i B\}_{i=1}^{r}$.*

Proof. Simply note that B acts linearly on the combined rowspace of the A_i, and this cannot increase the dimension. □

Proposition 9. *Given matrices $A_1, \ldots, A_r$ of size $m_1 \times n_1$ such that $\mathsf{crank}\{A_i\}_{i=1}^{r} \geq k_1$ and matrices $B_{i,j}$ for $1 \leq i \leq r$ and $1 \leq j \leq s$ of size $m_2 \times n_2$ such that $\mathsf{crank}\{B_{i,j}\}_{j=1}^{s} \geq k_2$ for each i we have,*

$$\mathsf{crank}\{A_i \otimes B_{i,j} : 1 \leq i \leq r, 1 \leq j \leq s\} \geq k_1 k_2.$$

Proof. Let U be an independent subset of $\cup_{i=1}^{r} \{u : u \text{ is a row in } A_i\}$ of size k_1 and for each $1 \leq i \leq r$ let V_i be an independent subset of $\cup_{j=1}^{s} \{v : v \text{ is a row in } B_{i,j}\}$ of size k_2. By the above tensor product bound we have that $\cup_{i=1}^{r} \{u \otimes v : u \in U, v \in V_i\}$ is a linearly independent set of size $k_1 k_2$. Now observe that if the rows of the A_i are regarded as coordinates of vectors u with respect to some basis $\{e_a\}_a$, and the rows of the $B_{i,j}$ are coordinates of vectors v with respect to some basis $\{f_b\}_b$, then the rows of the $A_i \otimes B_{i,j}$ represent the coordinates of the vectors $u \otimes v$ with respect to the basis $\{e_a \otimes f_b\}_{a,b}$. Thus the lower bound just proven is equivalent to a lower bound on the rank of the concatenated matrix, yielding the result. □

One last fact that will be necessary to sharpen our bounds at the end of the proof will be a tensor power trick. For this we will need to know that the repeated product of our Kakeya set is itself a Kakeya set - we prove that here.

Lemma 10. *Let K be a Kakeya set in R^n where $R = \mathbb{Z}/N\mathbb{Z}$ for a square-free integer $N = p_1 \dots p_r$. Then $K^m \subseteq R^{mn}$ is a Kakeya set in R^{mn}.*

Proof. Let $b \in \mathbb{P}R^{mn-1}$ be some direction and choose a representative $(b_1, \dots, b_m) \in (R^n)^m$ of this direction. Let $b_i^{(j)}$ denote the $\mathbb{F}_{p_j}^n$ component of b_i obtained through the Chinese remainder theorem. If all $b_i^{(j)}$ are non-zero, then every b_i corresponds to a direction in $\mathbb{P}R^{n-1}$ and the proof is straightforward, but some could be 0.
In the general case, for each i let $L_i \subseteq K$ be a line in some direction c_i that agrees with $b_i^{(j)}$ whenever it is non-zero. We claim that $L_1 \times \dots \times L_m$ contains a line in direction $(b_1, \dots, b_m)$. We have $L_i = \{a_i + tc_i : t \in R\}$ for some a_i. Let $L = \{(a_1, \dots, a_m) + t(b_1, \dots, b_m) : t \in R\}$, a line in direction b. Now, the set $L_1 \times \dots \times L_m$ contains all points of the form $(a_1 + t_1c_1, \dots, a_m + t_mc_m)$ where $t_i \in R$. Under the isomorphism given by the Chinese remainder theorem, choose t_i to be equal to t in the j^{th} entry when $b_i^{(j)}$ is non-zero, and 0 in the other entries. Then $t_ic_i = tb_i$ for each i, and we have $(a_1 + t_1c_1, \dots, a_m + t_mc_m) = (a_1, \dots, a_m) + t(b_1, \dots, b_m)$, so K^m contains a line in direction b and we are done. □

2.2 Proof of the Kakeya bound

We are now ready to develop the core part of the proof. We start with a definition and then construct the "decoding matrix" C_L^l.

Definition. Let m, n be natural numbers, and $U \subseteq \mathbb{F}^n$. We will consider vectors in $\mathbb{F}^{|U|\Delta_{n,m-1}}$ with entries indexed by (A, α) where A runs through U and α runs through the set of multiindices with $|\alpha| < m$. We define

$$\text{EVAL}_U^m : \mathbb{F}[X] \to \mathbb{F}^{|U|\Delta_{n,m-1}}$$

to be the linear map which sends a polynomial P to $(P^{(\alpha)}(A))_{(A,\alpha)}$. Here $P^{(\alpha)}$ is the α^{th} Hasse derivative.

Lemma 11. *Let $\mathbb{F}_q$ be a Finite Field, and let $l, n, m \in \mathbb{N}_0$ be such that $q|l$, $m = 2l - l/p$, and let $L \subseteq \mathbb{F}_q^n$ be a line in the direction $b \in \mathbb{P}\mathbb{F}_q^{n-1}$. Then we can construct a $\Delta_{n,l-1} \times q^n\Delta_{n,m-1}$ matrix C_L^l such that, for a homogeneous $P \in \mathbb{F}_q[X]$ of degree $lq - 1$ we have*

$$C_L^l \cdot EVAL_{\mathbb{F}_q^n}^m(P) = EVAL_b^l(P).$$

Moreover, following the notation from the previous definition, the only non-zero columns of C_L^l are the ones corresponding to (X, α) for which $X \in L$.

Proof. The bulk of the proof was contained in our proof of the Finite Field Kakeya bound. Indeed, we noted in the proof that for homogenous polynomials P of degree $lq - 1$ we have

$$\text{EVAL}_L^m(P) = 0 \Rightarrow \text{EVAL}_b^l(P) = 0$$

where $m = 2l - l/q$. Now, recall the basic fact from Linear Algebra that whenever A and B are linear maps from $\mathbb{F}^k$ to some other (possibly different) vector spaces, we have that if for each $v \in \mathbb{F}^k$, $Av = 0$ implies $Bv = 0$, then there exists C such that $CA = B$. This is straightforward, since this means that the kernel of A is a subset of the kernel of B, so the dimension of the range of B is

at most that of A, hence we have enough linearly independent vectors to construct a map C with $CA = B$.
These two facts together show that there is a matrix C' such that

$$C' \cdot \mathrm{EVAL}_L^m(P) = \mathrm{EVAL}_b^l(P).$$

We now add in zero columns to C' to correspond to (X, α) for $X \in \mathbb{F}_q^n \setminus L$, and we see that the resulting matrix C_L^l has the desired properties. □

With the decoding matrix now constructed, the proof now proceeds via induction and mostly makes use of simple counting arguments. Note that this proof also works for Kakeya sets in products of general Finite Fields (defined in the obvious way), but since we are more interested in $\mathbb{Z}/N\mathbb{Z} \cong \mathbb{F}_{p_1} \times \ldots \times \mathbb{F}_{p_r}$, we will use the notation p_i in the proof.

Proof of Theorem 2. We will use induction over r. For $r = 1$, this is simply Theorem 1, so suppose now that $r > 1$ and the result holds for $N_0 = p_2 \ldots p_r$. We will prove the result for $N = p_1 \ldots p_r$. Denote $R = \mathbb{Z}/N\mathbb{Z}$, $R_0 = \mathbb{Z}/N_0\mathbb{Z}$ and $p = p_1$ so that $R \cong \mathbb{F}_p \times R_0$. We will work over $\mathbb{F}_p$.
Let K be a Kakeya set in R^n. Every direction $b \in \mathbb{P}R^{n-1}$ is represented by $(b_1, b_2) \in \mathbb{P}\mathbb{F}_p^{n-1} \times \mathbb{P}R_0^{n-1}$. Through the Chinese remainder theorem, we see that a line $L \subseteq R^n$ in direction $b = (b_1, b_2)$ is a product of lines $L_1 \subseteq \mathbb{F}_p^n$ in direction b_1 and $L_2 \subseteq R_0^n$ in direction b_2.
Let I_L denote the indicator row vector of L, with entries $I_L(X)$ indexed by points of $X \in R^n$, and similarly for I_{L_1} and I_{L_2}. Identifying $X \in R^n$ with $(X_1, X_2) \in \mathbb{F}_p^n \times R_0^n$ by Chinese remainder theorem, we have $I_L(X) = I_{L_1}(X_1) I_{L_2}(X_2) = I_{L_1} \otimes I_{L_2}(X)$, the Kronecker product of I_{L_1} and I_{L_2}.
For each direction $b \in \mathbb{P}R^{n-1}$ we have at least one line in K in that direction. Pick one for each b and denote it by $L(b)$ contained in K. We may write it as a product $L_1(b) \times L_2(b)$ of lines in $\mathbb{F}_p^n$ and R_0^n in directions b_1 and b_2 respectively.
Now fix an l divisible by p, set $m = 2l - l/p$, and for a direction b consider the $\Delta_{n,l-1} \times p^n \Delta_{n,m-1}$ decoding matrix $C_{L_1(b)}^l$ over the field $\mathbb{F}_p$. We will show that

$$|K| \Delta_{n,m-1} \geq \mathsf{crank}\{C_{L_1(b)}^l \otimes I_{L_2(b)}\}_{b \in \mathbb{P}R^{n-1}}.$$

For each b, the columns in $C_{L_1(b)}^l$ are indexed by $(X, \alpha) \in \mathbb{F}_p^n \times \mathbb{N}_0^n$ with $|\alpha| < m$, hence the columns in $C_{L_1(b)}^l \otimes I_{L_2(b)}$ are indexed by $(X, \alpha) \in R^n \times \mathbb{N}_0^n$ with $|\alpha| < m$. The non-zero columns of $C_{L_1(b)}^l$ correspond to the points for which $X \in L_1(b)$, and so the non-zero columns in $C_{L_1(b)}^l \otimes I_{L_2(b)}$ correspond to points for which $X \in L(b) \subseteq K$. Hence the columns of the concatenated matrix are non-zero only if they correspond to (X, α) for which $X \in K$. For each such X there are $\Delta_{n,m-1}$ such columns, which gives the bound.
It remains to lower bound the crank of this set of matrices. For shorthand, let E be a matrix of size $p^n \Delta_{n,m-1} \times \delta_{n,lp-1}$ representing the linear map $\mathrm{EVAL}_{\mathbb{F}_p^n}^m$ restricted to the space of polynomials over $\mathbb{F}_p$ that are homogeneous of degree $lp - 1$. Given a direction $b_1 \in \mathbb{P}\mathbb{F}_p^{n-1}$, let D_{b_1} be the $\Delta_{n,l-1} \times \delta_{n,lp-1}$ matrix representing the linear map $\mathrm{EVAL}_{b_1}^l$ restricted to space of homogeneous polynomials of degree $lp - 1$.
For $b = (b_1, b_2) \in \mathbb{P}\mathbb{F}_p^{n-1} \times \mathbb{P}R_0^{n-1}$, we have $C_{L_1(b)}^l \cdot \mathrm{EVAL}_{\mathbb{F}_p^n}^m(P) = \mathrm{EVAL}_{b_1}^l(P)$ for any polynomial P homogeneous of degree $lp - 1$, that is, $C_{L_1(b)}^l E = D_{b_1}$. Let $I_{N_0^n}$ be an identity matrix of size $N_0^n \times N_0^n$. Then

$$\begin{aligned}
\mathsf{crank}\{C_{L_1(b)}^l \otimes I_{L_2(b)}\}_{b \in \mathbb{P}R^{n-1}} &\geq \mathsf{crank}\{(C_{L_1(b)}^l \otimes I_{L_2(b)})(E \otimes I_{N_0^n})\}_{b \in \mathbb{P}R^{n-1}} \\
&= \mathsf{crank}\{(C_{L_1(b)}^l E) \otimes I_{L_2(b)}\}_{b \in \mathbb{P}R^{n-1}} \\
&= \mathsf{crank}\{D_{b_1} \otimes I_{L_2(b_1,b_2)}\}_{b=(b_1,b_2) \in \mathbb{P}\mathbb{F}_p^{n-1} \times \mathbb{P}R_0^{n-1}}.
\end{aligned}$$

To lower bound this, note that by Proposition 9 it suffices to separately lower bound the crank of the collection of D_{b_1} as b_1 ranges over $\mathbb{PF}_p^{n-1}$ and for each b_1, to lower bound the crank of the collection of $I_{L_2(b_1,b_2)}$ as b_2 ranges over $\mathbb{P}R_0^{n-1}$.
First, we show that $\mathsf{crank}(\{D_{b_1}\}_{b_1 \in \mathbb{PF}_p^{n-1}}) \geq \delta_{n,lp-1}$. Let us consider the matrix D obtained by concatenating these matrices. Observe that this is precisely the matrix for the map $\mathrm{EVAL}^l_{\mathbb{PF}_p^{n-1}}$ (where we have chosen a representative for each direction), restricted to the space of homogeneous polynomials of degree $lp-1$. We claim that this map is injective, so that its rank is equal to the dimension of its domain, which is $\delta_{n,lp-1}$.
To see this, observe that if some homogeneous polynomial P lies in the kernel of this map, then all its Hasse derivatives of order less than l vanish over $\mathbb{PF}_p^{n-1}$. Since P is homogenous, so are its Hasse derivatives, hence P and its Hasse derivatives of order less than l vanish everywhere. By the extended Schwartz-Zippel lemma (more precisely, its corollary), as $(lp-1)p^{n-1} < lp^n$, we have $P = 0$, so the map is injective.
Next we show that for each $b_1 \in \mathbb{PF}_p^{n-1}$ we have

$$\mathsf{crank}\{I_{L_2(b_1,b_2)}\}_{b_2 \in \mathbb{P}R_0^{n-1}} \geq \frac{N_0^{n-1}}{\prod_{i=2}^{r}(2-1/p_i)^n}.$$

It is here that we use the inductive hypothesis - observe that for fixed b_1 the union of the $L_2(b_1, b_2)$ is a Kakeya set S in R_0^n. The crank of the set of indicator vectors is just the rank of the line matrix M_S over $\mathbb{F}_p$. By Proposition 5, the rank of M_S over any field is bounded below by $|S'|/|R_0|$, which by the induction hypothesis is at least the desired lower bound.
It follows from these bounds that

$$|K|\Delta_{n,m-1} \geq \frac{N_0^{n-1}}{\prod_{i=2}^{r}(2-1/p_i)^n}\delta_{n,lp-1}$$

For our particular choices, this is

$$|K|\binom{2l-l/p+n-1}{n} \geq \frac{N_0^{n-1}}{\prod_{i=2}^{r}(2-1/p_i)^n}\binom{lp+n-2}{n-1}.$$

Let l be a square multiple of p. We will use a tensor power trick. Apply the argument so far to the product of K with itself $\sqrt{l}$ times, which is a Kakeya set by Lemma 10. the above bound now becomes

$$|K|^{\sqrt{l}}\binom{2l-l/p+n\sqrt{l}-1}{n\sqrt{l}} \geq \frac{N_0^{n\sqrt{l}-1}}{\prod_{i=2}^{r}(2-1/p_i)^{n\sqrt{l}}}\binom{lp+n\sqrt{l}-2}{n\sqrt{l}-1}.$$

Rearranging the terms gives

$$|K|^{\sqrt{l}} \geq \frac{N_0^{n\sqrt{l}-1}}{\prod_{i=2}^{r}(2-1/p_i)^{n\sqrt{l}}}\frac{(lp+n\sqrt{l}-2)\dots(lp-1)}{(2l-l/p+n\sqrt{l}-1)\dots(2l-l/p-1)}n\sqrt{l}.$$

Take the $\sqrt{l}^{\text{th}}$ root on both sides and let $l \to \infty$ among the square multiples of p. By standard limits, one easily sees that $(n\sqrt{l})^{1/\sqrt{l}} \to 1$, and

$$\left(\frac{N_0^{n\sqrt{l}-1}}{\prod_{i=2}^{r} (2-1/p_i)^{n\sqrt{l}}} \right)^{1/\sqrt{l}} \to \frac{N_0^n}{\prod_{i=2}^{r} (2-1/p_i)^n}.$$

For the remaining fraction, divide the top and bottom by l, and note that for large l the numerator and denominator respectively are products of $n\sqrt{l}-1$ and $n\sqrt{l}$ terms that are arbitarily close to p and $(2-1/p)$, hence taking the $\sqrt{l}^{\text{th}}$ root and letting $l \to \infty$ gives

$$|K| \geq \frac{p^n N_0^n}{(2-1/p)\prod_{i=2}^{r} (2-1/p_i)^n}$$

which is the desired result. □

2.3 Additional comments

Note that in Theorem 2, we have not actually proven bounds of the form $|K| \geq C_{n,\varepsilon} N^{n-\varepsilon}$ (for each $\varepsilon > 0$). We shall take up such matters here for the sake of completeness, but it is only a matter of applying standard results from Analytic Number Theory at this stage.

An immediate corollary of Theorem 2 is that $|K| \geq N^n 2^{-rn}$, where r is the number of distinct prime factors. This is known to be $O(\log(N)/\log\log(N))$, hence $2^r = O(N^{1/\log\log(n)})$, and so $2^{-rn} \geq C N^{n/\log\log(N)}$. The estimates $|K| \geq C_{n,\varepsilon} N^{n-\varepsilon}$ for each $\varepsilon > 0$ follow immediately.

Let us include a brief proof of this fact. Let $\omega(N)$ denote the number of distinct prime factors of N. We have:

Lemma 12. *$\omega(N) = O(\log(N)/\log\log(N))$ as $N \to \infty$*

Proof. Firstly, note that by Stirling's formula (the ratio of $k!$ and $\sqrt{2\pi k}(k/e)^k$ approaches 1 as $k \to \infty$), we have $k! \geq (k/e)^k = e^{k\log(k)-k}$ for sufficiently large k. Take logs and set $k = \omega(N)$ to get $\omega(N)\log(\omega(N)) - \omega(N) \leq \log(\omega(N)!)$ (except for when $\omega(N)$ is small and the desired bound is trivial).

Now, we can write $N = p_1^{a_1} \dots p_r^{a_r}$ where $r = \omega(N)$ and p_i are prime numbers satisfying $p_1 < \dots < p_r$. Clearly $i < p_i$ for each i, so $r! < p_1 \dots p_r \leq N$, hence $\omega(N)! \leq N$.

Thus $\omega(N)\log(\omega(N)) - \omega(N) \leq \log(N)$ except for small $\omega(N)$. Rearranging gives $\omega(N) \leq \log(N)/(\omega(N)-1) \leq \log(N)$ for large $\omega(N)$. Alternatively, we could rearrange as $\omega(N)\log(\omega(N)) \leq \log(N) + \omega(N) \leq 2\log(N)$.

Now let $C > 2$ and suppose for a contradiction that there are infinitely many N with $\omega(N) \geq C\log(N)/\log\log(N)$. Then for these N we have

$$[C\log(N)/\log\log(N)] \cdot [\log(C\log(N)/\log\log(N))] \leq \omega(N)\log(\omega(N)) \leq 2\log(N).$$

Rearranging, we get

$$\log(C\log(N)/\log\log(N)) \leq (2/C)\log\log(N)$$

and taking exponentials gives

$$C\log(N)/\log\log(N) \leq (\log(N))^{2/C}.$$

Rearranging gives $C(\log(N))^{1-(2/C)}/\log\log(N) \leq 1$ for infinitely many N, but as $N \to \infty$ the left hand side goes to infinity, a contradiction. This completes the proof. □

References

[1] M. Dhar, Z. Dvir. *Proof of the Kakeya set conjecture over rings of integers modulo square-free N*. arXiv:2011.11225.

[2] Z. Dvir, S. Kopparty, S. Saraf, M. Sudan. *Extensions to the Method of Multiplicities, with applications to Kakeya Sets and Mergers*. SIAM Journal on Computing, 2013.

[3] J. Hickman, J. Wright. *The Fourier Restriction and Kakeya Problems over Rings of Integers Modulo N*. Discrete Analysis, 2018.