CHAPTER 4

Additive Energy and Balog-Szemerédi-Gowers

4.1. Introduction

Suppose that A and B are two finite sets in some ambient abelian group. We define the *normalised additive energy* $\omega_+(A,B)$ to be the number of quadruples $(a_1, b_1.a_2, b_2) \in A \times B \times A \times B$ with $a_1 + b_1 = a_2 + b_2$ divided by $|A|^{3/2}|B|^{3/2}$. We will often speak of the additive energy $\omega_+(A)$ of a single set A, by which we mean $\omega_+(A,A)$. Additive energy is intimately related to the sumset operation. One aspect of this relation is very easy to describe, and it asserts that small doubling implies large additive energy. More generally one has a "bilinear" version involving two different sets A and B.

Lemma 4.1 (Small sumset implies large additive energy). *Suppose that $\sigma[A,B] \leqslant K$. Then $\omega_+(A,B) \geqslant 1/K$. In particular if $\sigma[A] \leqslant K$ then $\omega_+(A) \geqslant 1/K$.*

Proof. For $x \in A+B$ write $r(x)$ for the number of pairs $(a,b) \in A \times B$ with $a+b=x$. The simplest of double-counting arguments gives $\sum r(x) = |A||B|$, and on the other hand $\sum r(x)^2$ is precisely the number of solutions to $a_1+b_1 = a_2+b_2$. It follows from the Cauchy-Schwarz inequality that

$$|A|^{3/2}|B|^{3/2}\omega_+(A,B) = \sum r(x)^2 = \sum_{x \in A+B} r(x)^2 \geqslant \frac{1}{|A+B|}\Big(\sum r(x)\Big)^2 = \frac{|A|^2|B|^2}{|A+B|} \geqslant \frac{1}{K}|A|^{3/2}|B|^{3/2},$$

as required. □

The most obvious converse to this fails: a set may have large additive energy without having large doubling. If $A_1 = \{1,\dots,n\}$ and $A_2 \subseteq \{n+1,\dots\}$ is an arbitrary set of size n , and if $A = A_1 \cup A_2$, then $\omega_+(A)$ is certainly at least 1/100, yet $\sigma[A]$ could well have size $\sim cn$ if A_2 has no particular additive structure.

In the preceding example a large piece of the set A, namely A_1, had considerable structure. One of the most useful results in additive combinatorics is the Balog-Szemerédi-Gowers theorem, which asserts that this situation is typical.

Theorem 4.1 (Balog-Szemerédi-Gowers). *Suppose that $\omega_+(A,B) \geqslant 1/K$. Then there are sets $A' \subseteq A$ and $B' \subseteq B$ with $|A'|/|A|, |B'|/|B| \geqslant cK^{-C}$ such that $\sigma[A',B'] \leqslant CK^C$.*

It follows from Ruzsa calculus that both $\sigma[A']$ and $\sigma[B']$ are at most CK^C as well. An immediate corollary of the theorem is the following assertion about one set rather than two, which was the version stated by both Balog-Szemerédi and Gowers.

COROLLARY 4.1 (Balog-Szemerédi-Gowers). *Suppose that $\omega_+(A) \geqslant 1/K$. Then there is a set $A' \subseteq A$ with $|A'|/|A| \geqslant cK^{-C}$ such that $\sigma[A'] \leqslant CK^C$.*

Though it is a more general statement, the proof of Theorem 4.1 is actually conceptually a little easier than a direct proof of Corollary 4.1. Our treatment follows Sudakov-Szemerédi-Vu [**?**], as written up in Tao and Vu's book.

4.2. A lemma about paths of length three

The heart of the proof of Theorem 4.1 is a lemma about bipartite graphs which is of interest in its own right.

LEMMA 4.2. *Suppose that G is a bipartite graph on vertex set $V \cup W$, where $|V| = |W| = n$, and with αn^2 edges all of which join a vertex in V to one in W. Then there are subsets $V' \subseteq V$ and $W' \subseteq W$ with $|V'|, |W'| \geqslant c\alpha^C n$ such that between every pair $v' \in V'$ and $w' \in W'$ there are at least $c\alpha^C n^2$ paths of length 3 in G.*

To prove this we must first establish a similar (but slightly weaker) lemma about paths of length two.

LEMMA 4.3. *Suppose that G is a bipartite graph on vertex set $V \cup W$, where $|V| = |W| = n$, and with αn^2 edges all of which join a vertex in V to one in W. Let $\eta > 0$ be a further parameter. Then there is a subset $V' \subseteq V$ with $|V'| \geqslant \alpha n/2$ such that between $(1-\eta)|V'|^2$ of the ordered pairs of points $(v_1, v_2) \in V' \times V'$ there are at least $\eta\alpha^2 n/2$ paths of length 2.*

Proof. If $x \in G$, write $N(x)$ for the neighbourhood of x in G, or in other words the set of vertices in G which are joined to x by an edge. Note that, since G is bipartite, $N(v) \subseteq W$ whenever $v \in V$ and $N(w) \subseteq V$ whenever $w \in W$.

Now by a double-counting argument, we have
$$\sum_{w \in W} \sum_{v \in V} 1_{vw \in E(G)} = \alpha n^2,$$
where $E(G)$ is of course the set of edges of G. Applying Cauchy-Schwarz to this gives
$$\sum_{w \in W} \sum_{v, v' \in V} 1_{vw \in E(G)} 1_{v'w \in E(G)} \geqslant \alpha^2 n^3,$$
or in other words
$$\mathbb{E}_{v, v' \in V} |N(v) \cap N(v')| \geqslant \alpha^2 n. \tag{4.1}$$
This constitutes the rather basic observation that, on average, pairs (v, v') have many common neighbours. Now say that two vertices v and v' are *extremely unfriendly* if $|N(v) \cap N(v')| < \eta\alpha^2 n/2$, or in other words if there are fewer than $\eta\alpha^2 n/2$

paths of length two between v and v'. Write $S \subseteq V \times V$ for the set of extremely unfriendly pairs. Manifestly, from (4.1), we have

$$\mathbb{E}_{v,v' \in V}(\eta - 1_{(v,v') \in S})|N(v) \cap N(v')| \geqslant \eta\alpha^2 n/2.$$

This may be rewritten as

$$\mathbb{E}_{v,v' \in V}(\eta - 1_{(v,v') \in S}) \sum_{w \in W} 1_{vw \in E(G)} 1_{v'w \in E(G)} \geqslant \eta\alpha^2 n/2.$$

Turning the sum over W into an expectation (by dividing by $|W| = n$) and swapping the order of summation, this implies that

$$\mathbb{E}_{w \in W}\mathbb{E}_{v,v' \in V}(\eta - 1_{(v,v') \in S})1_{v,v' \in N(w)} \geqslant \eta\alpha^2/2.$$

In particular there is a choice of w such that

$$\mathbb{E}_{v,v' \in V}(\eta - 1_{(v,v') \in S})1_{v,v' \in N(w)} \geqslant \eta\alpha^2/2.$$

Simply the fact that this expectation is greater than zero tells us that at most a proportion η of the pairs $v, v' \in N(w)$ are extremely unfriendly. Furthermore (ignoring the term involving S completely) we have

$$\mathbb{E}_{v,v' \in V} 1_{v,v' \in N(w)} \geqslant \alpha^2/2,$$

which implies that $|N(w)| \geqslant \alpha/\sqrt{2}$. Taking $V' := N(w)$, this proves the result. □

Remarks. This proof looks extremely slick at first sight. However when faced with the task of proving Lemm a4.3 it is not hard to develop the feeling that one must somehow select a very "connected" subset of V. The way we have done this is essentially by picking a random vertex $w \in W$, and taking V' to be the neighbourhood $N(w)$ of w in V, though this was easier to manage by using expectations rather than starting with "pick $w \in W$ uniformly at random and consider $N(w)$". This kind of technique seems to have been pioneered in this context by Gowers, and it is called "dependent random selection": one chooses something random (w in this case), then makes a deterministic choice based on it ($N(w)$).

It is not possible to guarantee that *all* pairs $v, v' \in V$ have many paths of length two between them; on the second example sheet you are asked to construct an example where this is not possible.

Let us turn now to the proof of Lemma 4.2. This is actually nowhere near as "clever" as the proof of the preceding lemma, but it is a little tedious.

Proof of Lemma 4.2. Delete all edges emanating from vertices in V with degree less than $\alpha n/2$; this causes the deletion of at most $\alpha n^2/2$ edges in total, so at least $\alpha n^2/2$ remain. From now on if we speak of an *edge* we mean one of these edges. Let $\eta > 0$ be a parameter to be chosen later. Using the preceding lemma, we may select a set $V' \subseteq V$ with $|V'| \geqslant \alpha n/4$ such that a proportion $1 - \eta$ of the pairs of vertices in V' have at least $\eta\alpha^2 n/8$ common neighbours in W.

All vertices in V' have degree 0 or else degree at least $\alpha n/2$, but it is conceivably the case that some do have degree 0. However if $\eta < 1/4$ then clearly no more than half of them do. Thus we may pass to a set $V'' \subseteq V'$, $|V''| \geqslant \alpha n/8$, such that every vertex in V'' has degree at least $\alpha n/2$ and still such that a proportion $1-\eta$ of the pairs of vertices in V'' have at least $\eta\alpha^2 n/8$ common neighbours in W.

Now let us focus on W. Look at all the edges from V'' into W: since each vertex in V'' has degree at least $\alpha n/2$, and $|V''| \geqslant \alpha n/8$, there are at least $\alpha^2 n^2/16$ of these. It follows that there is some set $W' \subseteq W$, $|W'| \geqslant \alpha^2 n/32$, such that each $w \in W'$ has at least $\alpha^2 n/32$ neighbours in V''.

Before concluding, let us jump back over to the other side and effect one final refinement of V''. Say that a vertex $v \in V''$ is *sociable* if there is a proportion at least $1-2\eta$ of the other vertices $v' \in V''$ are such that v and v' have at least $\eta\alpha^2 n/8$ common neighbours. Then at least half the vertices of V'' are sociable: call this set V''', so that $|V'''| \geqslant \alpha n/16$.

We now claim that for any $x \in V'''$ and $y \in W'$ there are many paths of length three between x and y (in the original graph G). Indeed by the choice of W' there must be at least $\alpha^2 n/32$ elements of V'' adjacent to y. There must also be at least $(1-2\eta)|V''|$ vertices of V'' which have at least $\eta\alpha^2 n/8$ common neighbours with x. Provided that $\alpha^2 n/32 \geqslant 3\eta|V''|$, which will be the case if $\eta \leqslant \alpha^2/96$, these two sets intersect in a set $\tilde{V} \subseteq V''$ of size at least $\eta|V''|$. Thus each element z of $\tilde{V}$ is adjacent to y, and has $\eta\alpha^2 n/8$ common neighbours with x. This clearly leads to at least $\eta^2\alpha^2|V''|n/8$ paths of length three between x and y.

The only constraints on η were that $\eta \leqslant 1/4$ and that $\eta \leqslant \alpha^2/96$. The latter is clearly the more severe constraint, so set $\eta := \alpha^2/96$. The lemma is then proved. □

4.3. Proof of the Balog-Szemerédi-Gowers theorem

In this section we deduce Theorem 4.1 from the graph-theoretic lemma of the previous section.

Suppose then that A, B are two sets in some abelian group G and that $\omega_+(A, B) \leqslant 1/K$. This means, of course, that there are at least $|A|^{3/2}|B|^{3/2}/K$ solutions to $a_1 - b_1 = a_2 - b_2$. Note that the number of solutions to this equation is at most $|A|^2|B|$, since once a_1, b_1 and a_2 are specified b_2 is uniquely determined. Therefore $|B| \leqslant K^2|A|$, and similarly $|A| \leqslant K^2|B|$.

By an argument almost identical to the one used in the proof of Theorem 2.3 there are many "popular differences" in $A - B$. Specifically, writing $s(x)$ for the number of pairs $(a, b) \in A \times B$ with $a - b = x$, there are at least $|A|^{1/2}|B|^{1/2}/2K$ values of x for which $s(x) \geqslant |A|^{1/2}|B|^{1/2}/2K$.

Define a bipartite graph G on vertex set $A \cup B$ by joining $a \in A$ to $b \in B$ by an edge if $a - b$ is a popular difference in the above sense. Then G has at least $|A||B|/4K^2$ edges. Let $n = \max(|A|, |B|)$, and "pad out" the smaller vertex class of G to obtain a new graph having n vertices in each class. Recalling that $K^{-2} \leqslant |A|/|B| \leqslant K^2$, this graph has at least αn^2 edges where $\alpha := 1/4K^4$.

Applying Lemma **??**, we may locate sets $A' \subseteq A$ and $B' \subseteq B$ with $|A'|/|A|, |B'|/|B| \geqslant cK^{-C}$ and such that for every $a' \in A'$ and $b' \in B'$ there are at least $c\alpha^C n^2$ paths of length 3 in G between a' and b'. This, of course, means that there at least $c\alpha^C n^2$ choices of $a'' \in A$ and $b'' \in B$ such that all three of $a' - b''$, $a'' - b''$ and $a'' - b'$ are popular.

Noting that $(a' - b') = (a' - b'') - (a'' - b'') + (a'' - b')$, it follows that $a' - b'$ can be written in at least $c\alpha^C n^2$ ways as $x - y + z$, where x, y and z are popular differences. These are genuinely distinct representations, since it is easy to recover a'' and b'' from knowledge of a', b', x, y and z. However the number of popular differences is bounded *above* by $2K|A|^{1/2}|B|^{1/2} \leqslant C\alpha^{-C}n$, as can be seen by simply double-counting pairs $(a, b) \in A \times B$. It follows that

$$|A' - B'| \cdot c\alpha^C n^2 \leqslant (C\alpha^{-C}n)^3,$$

which of course implies that $|A' - B'| \leqslant C\alpha^{-C}n \leqslant C'K^{C'}n$. In view of the lower bounds for $|A'|$ and $|B'|$ already obtained, this clearly implies that $|A' - B'| \leqslant CK^C|A'|^{1/2}|B'|^{1/2}$ and hence by Ruzsa calculus that $\sigma[A', B'] \leqslant CK^C$, as claimed. □