Minkowski Complexity of Sets: An Easy Lower Bound

Stasys Jukna

Abstract. The Minkowski complexity of a finite set of vectors is the minimum number of set-theoretic union and Minkowski sum operations needed to create this set when starting from single-element sets, each containing only one vector. We give an amazingly simple proof of a general lower bound on this complexity.

1. INTRODUCTION. Let $\mathbb{N} = \{0, 1, 2, \ldots\}$. The *Minkowski sum* or *sumset* of two sets X and Y of vectors in $\mathbb{N}^n$ is the set

$$X + Y = \{x + y : x \in X \text{ and } y \in Y\}$$

of vectors in $\mathbb{N}^n$, where $x + y = (x_1 + y_1, \ldots, x_n + y_n)$ is the componentwise sum of vectors x and y. That is, we add every vector of X to every vector of Y.

Given a finite set $A \subset \mathbb{N}^n$ of vectors, how difficult is it to create this set when starting from single-element sets $\{x\}$ with $x \in \mathbb{N}^n$, and using set-theoretic union $X \cup Y$ and sumset $X + Y$ operations?

Every such process for creating A is a sequence $A_1, \ldots, A_l$ of sets of vectors such that $A_l = A$ and every set A_i is either the union $A_i = X \cup Y$ or the sumset $A_i = X + Y$ of two sets, each of which is either a single-element set $\{x\}$ with $x \in \mathbb{N}^n$ or some previously created set A_j with $j < i$. We call the minimal number l of steps needed to create A using such a process the *Minkowski complexity* of A, and denote it by $L(A)$.

Example 1. It is clear that $L(A) \leq |A| - 1$ holds for any set A, where $|A|$ is the cardinality of A: just take the union of all sets $\{x\}$ with $x \in A$. However, the usage of sumset operations can dramatically reduce the total number of operations. To give a simple example, define the *weight* of a vector as the sum of its entries, and let A be the set of all vectors in $\mathbb{N}^n$ of weight $N = 2^n$. That is, A consists of all nonnegative integer solutions of the equation $x_1 + x_2 + \cdots + x_n = N$. It is well known (and easy to show) that this set has about $(N/n)^n$ vectors, more exactly, $|A| = \binom{N+n-1}{n-1}$ (see, for example, [**1**, Theorem 2.2]). On the other hand, the Minkowski complexity of this set is very small: we can create A by using only $2n - 1$ operations. First, we use $n - 1$ union operations to create the set $A_0 = \{e_1, \ldots, e_n\}$, where e_i is the 0-1 unit vector of length n with exactly one 1 in the ith position. Then we use n sumset operations to create sets $A_1 = A_0 + A_0$, $A_2 = A_1 + A_1, \ldots, A_n = A_{n-1} + A_{n-1}$. Since each application of a sumset operation duplicates the weight, each set A_i is the set of all vectors of weight 2^i; hence, $A_n = A$, as desired.

Motivation. Minkowski complexity of sets of vectors is related to the algebraic complexity of multivariate polynomials $f(z_1, \ldots, z_n) = \sum_{a \in A_f} c_a \prod_{i=1}^{n} z_i^{a_i}$, where $A_f \subset \mathbb{N}^n$ is a finite set of exponent vectors, and all coefficients c_a are positive integers. The monotone algebraic complexity of a polynomial is the minimum number of sum and product operations needed to compute this polynomial when starting from variables

http://dx.doi.org/10.4169/amer.math.monthly.124.8.749
MSC: Primary 68Q17

$z_1, \dots, z_n$ and nonnegative integers; the word "monotone" stands here to stress that subtraction operations are not allowed.

There is a natural homomorphism from the semiring of such polynomials to the semiring $(2^{\mathbb{N}^n}, \cup, +)$ of finite subsets of vectors that maps every polynomial f to the set A_f of its exponent vectors. In particular, every single variable z_i is mapped to $A_{z_i} = \{e_i\}$. That this is indeed a homomorphism follows from easily verifiable equalities $A_{f+h} = A_f \cup A_h$ and $A_{f \cdot h} = A_f + A_h$, the latter sum being the sumset of A_f and A_h. Thus, the Minkowski complexity of the set of exponent vectors of a polynomial is a lower bound on the monotone algebraic complexity of this polynomial. In fact, *all* known lower bounds on the monotone algebraic complexity of polynomials, including those in [**3, 5**], were obtained by proving lower bounds on the Minkowski complexity of their sets of exponent vectors.

2. SIDON SETS. Which sets A have high Minkowski complexity? It turns out that such sets are the well-known Sidon sets. These are the sets $A \subseteq \mathbb{N}^n$ with the following property: if we know the sum of two vectors of A, then we know which vectors were added. More formally, A is a *Sidon set* if for any vectors a, b, c, d in A, $a + b = c + d$ implies $\{c, d\} = \{a, b\}$. For example, in the case $n = 1$, both sets $A = \{1, 2, 5, 7\}$ and $B = \{1, 3, 6, 7\}$ are Sidon sets, but $C = \{1, 2, 4, 5, 7\}$ is not a Sidon set because, for example, $2 + 4 = 1 + 5$. The term "Sidon set" was coined by Erdős and Turán [**2**] in honor of Simon Sidon who introduced these sets in order to solve a problem in harmonic analysis.

Example 2. The following *combinatorial* construction of large Sidon sets $A \subset \{0, 1\}^n$ is due to Schnorr [**5**]. Let $n = \binom{m}{2}$, and let A be the set of all $|A| = \binom{m}{k}$ characteristic 0-1 vectors of k-cliques, viewed as sets of their edges in a complete graph K_m on m nodes; a *k-clique* is obtained by taking any k-element subset of nodes and drawing edges between all $\binom{k}{2}$ pairs of these nodes. It is easy to verify that no union of two k-cliques can contain some third k-clique. Indeed, the latter clique must then have a node u not in the first clique and a node v not in the second clique. If $u = v$ then the node u is not covered, and if $u \neq v$ then the edge $\{u, v\}$ is not covered by the first two cliques, a contradiction. Thus, the set A has an even stronger property (than just being a Sidon set): if $a + b \geq c$ for some vectors $a, b, c \in A$, then $c \in \{a, b\}$. Such sets are also known as *cover-free* sets.

Example 3. The following *algebraic* construction of large Sidon sets is due to Lindström [**4**]. Let $A \subset \{0, 1\}^{2n}$ be the set of all $|A| = 2^n$ vectors (x, x^3) with x in $\{0, 1\}^n$, where we view vectors x as elements of GF(2^n)—that is, as coefficient vectors of polynomials of degree at most $n - 1$ over GF(2)—when raising them to a power. To show that A is a Sidon set, fix any two vectors $c, d \in \{0, 1\}^n$, and consider the equation $(x, x^3) + (y, y^3) = (c, c^3) + (d, d^3)$. It is enough to show that this equation has at most one unordered pair $\{x, y\}$ of 0-1 solutions over the semigroup $(\mathbb{N}^{2n}, +)$. If $c = d$, then there is only one solution $x = y = c$. So, assume that $c \neq d$. It is enough to show that then the equation cannot have more than one solution $\{x, y\}$ even over the field GF(2^{2n}). The equation is equivalent to the system of two equations $x + y = a$ and $x^3 + y^3 = b$ with $a = c + d \neq 0$ and $b = c^3 + d^3$. Since we are working over a field of characteristic 2, we have $-x = x$ and $3x = x$. So, the identity $3xy(x + y) = (x + y)^3 - (x^3 + y^3) = a^3 - b$ turns into $xy = a^2 + b/a$. Thus, x and y must satisfy $x + y = a \neq 0$ and $xy = a^2 + b/a$. By Vieta's formulas, x and y are then the solutions of the quadratic equation $aX^2 + a^2X + (a^3 + b) = 0$, and there can be only one pair of them.

3. THE THEOREM. Our goal is to present an amazingly simple proof of a general lower bound on the Minkowski complexity, implying that Sidon sets have almost maximal complexity.

To do this, we associate with every finite set $X \subset \mathbb{N}^n$ of vectors its nonnegative real *cost* $\mu(X)$. Such a cost measure is *legal* if the following three natural conditions are fulfilled: $\mu(\{x\}) \leq 1$ for every vector $x \in \mathbb{N}^n$, $\mu(X \cup Y) \leq \mu(X) + \mu(Y)$, and $\mu(X + Y) \leq \mu(X) \cdot \mu(Y)$. For example, a trivial cost measure $\mu(X) = |X|$, the cardinality of X, is legal.

A set A is *k-free* ($k \geq 1$) with respect to a given cost measure if $X + Y \subseteq A$ implies that at least one of the sets X and Y must have cost at most k. Note that there is no restriction on the cost of the other set: it may be arbitrarily expensive. The restriction is that they cannot *both* be expensive.

Theorem. *If a set $A \subset \mathbb{N}^n$ is k-free with respect to some cost measure μ, then*

$$L(A) \geq \mu(A)/2k^3.$$

Corollary. *If $A \subset \mathbb{N}^n$ is a Sidon set, then $L(A) \geq |A|/2$.*

Proof. Let $A \subseteq \mathbb{N}^n$ be a Sidon set. It is enough to show that then A must be 1-free with respect to a trivial cost measure $\mu(X) = |X|$, the cardinality of X.

Indeed, were A not 1-free, then we would have $\{x, x'\} + \{y, y'\} \subseteq A$ for some vectors $x \neq x'$ and $y \neq y'$. The sum $a + b$ of the two vectors $a = x + y$ and $b = x' + y'$ of A is then equal to the sum $c + d$ of the two vectors $c = x + y'$ and $d = x' + y$ of A. Since A is a Sidon set, at least one of $x + y' = x + y$ or $x + y' = x' + y'$ must hold, contradicting that $x \neq x'$ and $y \neq y'$. ■

This latter result is not quite new: by improving a classical result of Schnorr [**5**, Theorem 3.2], Gashkov and Sergeev [**3**, Theorem 1] have shown that Sidon sets A have Minkowski complexity at least $|A|$ when *only* $n + 1$ single-element sets $\{0\}, \{e_1\}, \dots, \{e_n\}$ are allowed as initial sets; here $0 \in \mathbb{N}^n$ is the all-0 vector. Recall that we allow *any* sets $\{x\}$ with $x \in \mathbb{N}^n$ to start with, and hence, $L(A) \leq |A| - 1$ is a trivial upper bound in our (more general) setting.

The proofs in [**3, 5**] are ingenious but rather technical. Our contribution is the simplicity of the proof.

4. SUMSET NETWORKS. It will be convenient to view legal processes of creating a given set as graphs. Recall that a *digraph* (directed graph) is a pair $G = (V, E)$, where V is a finite set of *nodes*, and $E \subseteq V \times V$ is a set of *edges* $e = (u, v)$; node u is the *tail*, and node v is the *head* of e. We say that the edge e *leaves* its tail u and *enters* its head v. The *indegree* (resp., *outdegree*) of a node is the number of edges entering (resp., leaving) this node. A graph is *acyclic* if there are no closed walks in it.

Define a *sumset network* to be an acyclic digraph with one zero outdegree node (the *output* node), and some number of zero indegree nodes (*source* nodes), each holding some single-element set $\{x\}$ with $x \in \mathbb{N}^n$. Every other node, a *gate*,[1] has indegree two, and performs either the union or the sumset operation on its two inputs.

At each node v of such a network, some set $X_v \subseteq \mathbb{N}^n$ of vectors is created in a natural way. If v is a source node holding a single-element set $\{x\}$, then $X_v = \{x\}$. If

[1] The term "gate" comes from electrical engineering, and is only used to stress that a node has its associated operation.

v is a gate entered by nodes u and w, then $X_v = X_u \cup X_w$ if v is a union gate, and $X_v = X_u + X_w$ if v is a sumset gate. The set created by the entire network is the set created at the output gate. Thus, the Minkowski complexity of a given set $A \subset \mathbb{N}^n$ is exactly the minimum number of gates in a sumset network producing this set. Note that we only count gates: source nodes are given "for free."

Fix a sumset network, and let $A \subset \mathbb{N}^n$ be the set of vectors created by this network. We associate with every gate v the following three sets of vectors:

- X_v is the set of vectors created at gate v (as defined above);
- $Y_v = \{y \in \mathbb{N}^n : x + y \in A \text{ for all } x \in X_v\}$ is the *residue* of X_v (within A);
- $Z_v = X_v + Y_v$ is the *content* of gate v.

Note that X_v need not lie in A, but at least one of its translates $X_v + y$ must already lie in A; the residue Y_v collects all such vectors y. For example, if v is the output gate, then $X_v = A$ and $Y_v = \{0\}$. If v is a source node holding a set $\{x\}$ with $x \in \mathbb{N}^n$, then $X_v = \{x\}$ and Y_v consists of all vectors $y \in \mathbb{N}^n$ such that $x + y \in A$.

We define the *content* Z_e of an edge $e = (u, v)$ to be the content $Z_v = X_v + Y_v$ of its head v if v is a sumset gate. But if v is a union gate, then we define the content Z_e of edge e as the sumset $X_u + Y_v$. The reason for the asymmetry in this definition is explained by the following *content propagation property*. Let v be a gate and $a \in Z_v$ be a vector in its content.

- If v is a union gate, then vector a belongs to the content of *at least one* edge entering v, as well as to the content of the tail of this edge.
- If v is a sumset gate, then vector a belongs to the contents of *both* edges entering v, as well as to the contents of the tails of these edges.

Proof. The gate v is entered by edges from some two gates u and w. First let v be a union gate. Since vector a belongs to the content $X_v + Y_v = (X_u \cup X_w) + Y_v$ of gate v, the vector a must belong to at least one of the contents $X_u + Y_v$ or $X_w + Y_v$ of the edges entering gate v. Assume without loss of generality. that a belongs to $X_u + Y_v$. Since in the case of a union gate we have that $Y_v = Y_u \cap Y_w$, vector a belongs to the content $X_u + Y_u$ of the tail u of edge (u, v) as well.

Now let v be a sumset gate. In this case, the contents of both edges (u, v) and (w, v) coincide with the content $X_v + Y_v$ of gate v. So, vector a belongs to the contents of both of these edges. To show that a must also belong to the contents of the tails u and w of these edges, we first show that we have an inclusion $X_w + Y_v \subseteq Y_u$. Indeed, were there two vectors $x \in X_w$ and $y \in Y_v$ for which $x + y$ does not belong to Y_u, then, by the definition of the residue Y_u, there should be an $x' \in X_u$ for which $x' + (x + y)$ does not belong to A. But this is impossible because $x' + x$ belongs to the set $X_v = X_u + X_w$ created at the gate v, and y belongs to the residue Y_v of X_v. Now, the inclusion yields $X_v + Y_v = (X_u + X_w) + Y_v = X_u + (X_w + Y_v) \subseteq X_u + Y_u$, and similarly, $X_v + Y_v \subseteq X_w + Y_w$. Thus, every vector in the content $X_v + Y_v$ of v (including our vector a) belongs to the contents of both input gates u and w, as desired. ∎

5. PROOF OF THE THEOREM. Let $A \subset \mathbb{N}^n$ be some k-free set of vectors. That is, it contains no sumset $X + Y$ such that *both* X and Y have cost larger than k. Fix an arbitrary sumset network producing A. Our goal is to show that the number of gates in this network must be at least the cost $\mu(A)$ of A divided by $2k^3$. Call an edge e in the network *cheap* if the cost $\mu(Z_e)$ of its content Z_e is at most k^3. We claim that *every vector of A belongs to the content of at least one cheap edge.*

Given this claim, the theorem follows easily. Indeed, if E is the set of all cheap edges in the network, then the claim implies that $A \subseteq \bigcup_{e \in E} Z_e$. Since $\mu(Z_e) \leq k^3$ holds for every edge $e \in E$, and since the cost of a union of two sets is at most the sum of their costs, this implies $\mu(A) \leq \sum_{e \in E} \mu(Z_e) \leq k^3|E|$. So, there must be at least $|E| \geq \mu(A)/k^3$ edges in the network. Since every gate (nonsource node) in the network is entered by exactly two edges, and every edge must enter some gate, the total number of edges in the network is twice the total number t of gates. Hence, $t \geq |E|/2 \geq \mu(A)/2k^3$, as claimed.

So, it remains to prove the claim. Call a node u in the network *cheap* if $\mu(X_u) \leq k$, and *expensive* otherwise. We can assume that the output gate is expensive because otherwise the set A created by the network would be also cheap, and there would be nothing to prove.

Fix a vector $a \in A$. Start at the output gate of the network, and construct a path in the underlying directed acyclic graph by going backwards and using the following rule, where v is the last previously reached node.

(a) If v is a union gate, then go to an input whose content contains a.

(b) If v is a sumset gate, then go to either of the two inputs if they are both expensive or are both cheap, and go to the expensive input if the second input is cheap.

Every source node has cost at most $1 \leq k$, and hence, is cheap. Since the output gate is expensive, we will eventually reach some source node. The content propagation property ensures that the vector a must belong to the contents of *all* edges along the corresponding source-to-output path. Since the first node of this path (a source node) is cheap, and the last one (output gate) is expensive, there must be an edge $e = (u, v)$ such that $\mu(X_u) \leq k$ but $\mu(X_v) > k$. It remains to show that this edge e must be cheap. Since $X_v + Y_v \subseteq A$ and $\mu(X_v) > k$, the k-freeness of A implies that $\mu(Y_v) \leq k$ must hold.

If v is a union gate, then $\mu(Z_e) = \mu(X_u + Y_v) \leq \mu(X_u) \cdot \mu(Y_v) \leq k^2 \leq k^3$.

If v is a sumset gate, and w is its second input, then the content of e is the Minkowski sum $Z_e = X_v + Y_v = X_u + X_w + Y_v$ of three sets. Since the node u is cheap, rule (b) in the construction of the path implies that the second node w entering v must have also been cheap, that is, $\mu(X_w) \leq k$ must hold. We thus have $\mu(Z_e) \leq \mu(X_u) \cdot \mu(X_w) \cdot \mu(Y_v) \leq k^3$, as desired. ■

ACKNOWLEDGMENTS. I thank Georg Schnitger and Igor Sergeev for numerous inspiring discussions. I am also grateful to the referee for useful comments on presentation. This research was supported by DFG grant SCHN 503/6-1.

REFERENCES

1. M. Bona, *Introduction to Enumerative and Analytic Combinatorics*. Second ed. CRC Press, Boca Raton, FL, 2015.
2. P. Erdős, P. Turán, On a problem of Sidon in additive number theory, and on some related problems, *J. Lond. Math. Soc.* **16** (1941) 212–215.
3. S. Gashkov, I. Sergeev, On a method for deriving lower bounds for the complexity of monotone arithmetic circuits computing real polynomials, *Mat. Sb.* **203** (2012) 1411–1147.
4. B. Lindström, Determination of two vectors from the sum, *J. Combin. Theory Ser. B* **6** (1969) 402–407.
5. C. Schnorr, A lower bound on the number of additions in monotone computations, *Theoret. Comput. Sci.* **2** (1976) 305–315.

Institute of Computer Science, Goethe University Frankfurt, Frankfurt am Main, Germany
Affiliated with Institute of Mathematics and Informatics, Vilnius University, Vilnius, Lithuania
jukna@thi.informatik.uni-frankfurt.de