Irreducible Subcube Partitions

Yuval Filmus, Edward Hirsch, Sascha Kurz, Ferdinand Ihringer,
Artur Riazanov, Alexander Smal, Marc Vinyals

December 30, 2022

Abstract

A *subcube partition* is a partition of the Boolean cube $\{0,1\}^n$ into subcubes. A subcube partition is *irreducible* if the only sub-partitions whose union is a subcube are singletons and the entire partition. A subcube partition is *tight* if it "mentions" all coordinates.

We study extremal properties of tight irreducible subcube partitions: minimal size, minimal weight, maximal number of points, maximal size, and maximal minimum dimension. We also consider the existence of *homogeneous* tight irreducible subcube partitions, in which all subcubes have the same dimensions. We additionally study subcube partitions of $\{0,\ldots,q-1\}^n$, and partitions of $\mathbb{F}_2^n$ into affine subspaces, in both cases focusing on the minimal size.

Our constructions and computer experiments lead to several conjectures on the extremal values of the aforementioned properties.

1 Introduction

A *subcube partition* is a partition of the cube $\{0,1\}^n$ into subcubes, that is, into sets of the form

$$\{x \in \{0,1\}^n : x_{i_1} = b_1, \ldots, x_{i_d} = b_d\}.$$

Here is an example of a subcube partition of *length* $n = 3$:

$$\begin{aligned} S_3 &= \{000\}, \{111\}, \{001, 101\}, \{100, 110\}, \{010, 011\} \\ &= 000, 111, {*}01, 1{*}0, 01{*}. \end{aligned}$$

We will usually express our subcubes as strings in $\{0,1,*\}^n$, in which stars stand for unconstrained coordinates.

A subcube partition is *reducible* if it has a proper subset, consisting of more than one subcube, whose union is a subcube. For example,

$$0{*}, 10, 11$$

is reducible since $10 \cup 11 = 1{*}$. In contrast, S_3 is *irreducible*.

A subcube partition is *tight* if it *mentions* all coordinates, that is, if for every $i \in [n]$, some subcube constrains x_i. Both subcube partitions above are tight, but the subcube partition $0{*}, 1{*}$ is not, since the second coordinate is not mentioned.

Peitl and Szeider [PS22] enumerated all tight irreducible subcube partitions for $n = 3, 4$, and counted the number of nonisomorphic subcube partitions with small *size* (number of subcubes) for $n = 5, 6, 7$. They ask whether there are infinitely many tight irreducible subcube partitions. In this work, we answer this question in the affirmative, giving many constructions of tight irreducible subcube partitions.

The work of Peitl and Szeider raises many natural questions, such as:

- How to determine whether a subcube partition is irreducible?

- What is the minimal size of a tight irreducible subcube partition of length n?
  (This question only makes sense if we impose tightness.)
- What is the maximal size of an irreducible subcube partition of length n?
- Do there exist irreducible subcube partitions in which all subcubes have the same dimension?
  (We call such subcube partitions *homogeneous*.)

We address these questions in Section 2. We describe an efficient algorithm for testing whether a subcube partition is irreducible in Section 2.1, and give an infinite sequence of irreducible formulas in Section 2.2.

We conjecture that the minimal size of a tight irreducible subcube partition of length n is $2n-1$. We give a matching construction in Section 2.3, and optimize its Hamming weight in Section 2.4 (this will be useful later on).

We conjecture that the maximal size of an irreducible subcube partition of length $n \geq 5$ is $\frac{5}{8}2^n$. We present constructions of such subcube partitions for small n in Section 2.6, where we also give a nontrivial upper bound. Our constructions involve 2^{n-2} points (0-dimensional subcubes) and $3{\cdot}2^{n-3}$ edges (1-dimension subcubes). We conjecture that 2^{n-2} is the maximal number of points in an irreducible subcube partition of size n. A matching construction appears in Section 2.5.

We present subcube partitions in which all subcubes have linear dimension in Section 2.7. We close off the section with a discussion of homogeneous subcube partitions in Section 2.8, where we describe several constructions, and determine all lengths n and codimensions $k \leq 4$ for which there exists a tight irreducible subcube partitions of length n whose subcubes have dimension $n-k$.

Section 3 studies subcube partitions of $[q]^n$ for $q \geq 3$. We show how to construct irreducible subcube partitions of $[q]^n$ from irreducible subcube partitions of $\{0,1\}^n$ in Section 3.1, and use this to construct tight irreducible subcube partitions of length n and size $(n-1)q(q-1)+1$ in Section 3.2.1; this uses the subcube partitions of Section 2.4. We conjecture that $(n-1)q(q-1)+1$ is the minimum size of a tight irreducible subcube partition, and prove this for $n=3$ in Section 3.2.2. We close by showing, in Section 3.2.3, that the minimal size of a tight *minimal cover* in this setting is $(q-1)n+1$.

Finally, Section 4 briefly studies the linear analog of subcube partitions, in which subcubes are replaced by affine subspaces. We show how to construct irreducible affine vector space partitions from irreducible subcube partitions in Section 4.1, and use this to construct tight irreducible affine subspace partitions of length n and size roughly $\frac{3}{2}n$ in Section 4.2. We discuss irreducible affine vector space partitions in more detail in the companion work [BFIK22].

Background

Subcube partitions appear, under various names, in theoretical computer science, as an abstraction of the salient properties of decision trees, and elsewhere. Some examples include Iwama [Iwa87, Iwa89] (as certain *independent sets of clauses*), Brandman, Orlitsky and Hennessy [BOH90] (as *nonoverlapping covers*), Berger, Felzanbaum and Fraenkel [BFF90] (as *disjoint tautologies*), Davydov and Davydova [DD98] (as *dividing formulas*), Friedgut, Kahn and Wigderson [FKW02] (as *subcube partitions*), Kullmann [Kul04] (as *unsatisfiable hitting clause-sets*), Kisielewicz [Kis20] (as realizations of *cube tiling codes*). There are also *orthogonal DNFs* [CH11], also known as *disjoint DNFs* [GK13], which are systems of disjoint subcubes which do not necessarily cover the entire cube. (For the relation between decision trees and subcube partitions, see Göös, Pitassi and Watson [GPW18].)

Irreducible subcube partitions appear in work of Kullmann and Zhao [KZ16] (as *clause-reducibility*), inspired by similar notions in the context of disjoint covering systems of residue classes [Kor84, BFF90] and motivated by applications to the study of CNFs.

Peitl and Szeider [PS22] enumerate all tight irreducible subcube partitions for $n = 3, 4$, and determine the minimal size of a regular irreducible subcube partition for $n = 5, 6, 7$. Instead of tightness, they use a different notion, *regularity*, which is equivalent to tightness for irreducible subcube partitions when $n \geq 3$. Regularity was introduced by Kullmann and Zhao [KZ13] under the name *nonsingularity*, and is defined in Section 2.3.1.

Acknowledgements We thank Jan Johannsen, Ilario Bonacina, Oliver Kullmann, and Stefan Szeider for introducing us to the topic; Tomáš Peitl and Stefan Szeider for sharing with us the results of the computer search reported in [PS22]; and Zachary Chase for helpful discussions.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 802020-ERC-HARMONIC. Ferdinand Ihringer is supported by a postdoctoral fellowship of the Research Foundation – Flanders (FWO).

2 Subcube partitions

We start with a quick recap of the relevant definitions.

Definition 2.1 (Subcube partition)**.** A *subcube partition* of length n is a partition of $\{0,1\}^n$ into *subcubes*, which are sets of the form

$$\{x \in \{0,1\}^n : x_{i_1} = b_1, \ldots, x_{i_d} = b_d\}.$$

The parameter d is the *codimension* of the subcube, and $n-d$ is its *dimension*. A subcube of dimension 0 is called a *point*, and a subcube of dimension 1 is called an *edge*.

The *size* of a subcube partition is the number of subcubes.

We identify subcubes with words over $\{0,1,*\}$. For example, $01*$ stands for the subcube $\{(0,1,0),(0,1,1)\}$. We index the symbols in a word w of length n by $[n] = \{1,\ldots,n\}$. If $b \in \{0,1\}$, we use $\bar{b}$ to denote $1-b$.

Definition 2.2 (Reducibility)**.** A subcube partition F is *reducible* if there exists a subset $G \subset F$, with $1 < |G| < |F|$, such that the union of the subcubes in G is itself a subcube.

A subcube partition is *irreducible* if it is not reducible.

Definition 2.3 (Tightness)**.** A subcube s *mentions* a coordinate $i \in [n]$ if $s_i \neq *$.

A subcube partition F of length n is *tight* if for every $i \in [n]$, some subcube in F mentions i.

It is coNP-complete to determine whether a given collection of subcubes covers $\{0,1\}^n$ (this problem is just SAT in disguise). In contrast, it is easy to test whether a given collection of subcubes is a partition, as first observed by Iwama [Iwa89].

Definition 2.4 (Conflicting subcubes)**.** Two subcubes s,t of the same length are said to *conflict* if there is a coordinate $i \in [n]$ such that $s_i, t_i \neq *$ and $s_i \neq t_i$.

Lemma 2.5. *Two subcubes are disjoint if and only if they conflict.*

Lemma 2.6. *A collection F of disjoint subcubes of length n is a subcube partition if and only if*

$$\sum_{s \in F} 2^{-\operatorname{codim}(s)} = 1.$$

Similarly, it is easy to check whether a given subcube partition is tight. In contrast, checking whether a subcube partition is irreducible using the definition takes exponential time. We present an efficient algorithm for testing irreducibility in Section 2.1.

Following that, we give many examples of irreducible subcube partitions, starting with Section 2.2, which describes a one-off construction. In Sections 2.3 to 2.7 we describe irreducible subcube partitions which conjecturally optimize various parameters. Section 2.8 closes with a discussion of irreducible subcube partitions in which all subcubes have the same dimension.

2.1 Testing irreducibility

In this section we give a polynomial time algorithm that checks whether a given subcube partition F is reducible, and if so, identifies a subset $G \subset F$, with $1 < |G| < |F|$, whose union is a subcube.

The idea behind the algorithm is quite simple. Suppose that F were reducible, say via the subset G. If $s, t \in G$ then $\bigcup G$ must contain the *join* $s \vee t$ of s, t, which is the smallest subcube containing both s and t, given explicitly by

$$(s \vee t)_i = \begin{cases} b & \text{if } s_i = t_i = b \in \{0,1\}, \\ * & \text{otherwise.} \end{cases}$$

If $u \in F$ intersects $s \vee t$ (a condition we can check using Lemma 2.5) then G must contain u, and so $\bigcup G$ must contain $s \vee t \vee u$. Continuing in this way, we are able to recover G (or a subset of G whose union is also a subcube). The corresponding algorithm appears as Algorithm 1.

Algorithm 1 Algorithm for checking whether a subcube partition is irreducible

Input: Subcube partition $F = \{s_1, \ldots, s_m\}$
for $1 \le i < j \le m$ **do**
  $G \leftarrow \{s_i, s_j\}$
  while some $s_k \notin G$ intersects $\bigvee G$ **do**
    $G \leftarrow G \cup \{s_k\}$
  end while
  if $G \neq F$ **then**
    return Reducible: $\bigcup G$ is a subcube
  end if
end for
return Irreducible

Theorem 2.7. *Algorithm 1 runs in polynomial time, and its output is correct.*

Proof. We start by showing that the algorithm runs in polynomial time. The outer **for** loop runs $O(m^2)$ times, and the inner **while** loop runs at most m times. Each basic operation can be implemented in polynomial time, and so the entire algorithm runs in polynomial time.

Suppose first that the algorithm outputs "reducible". By construction, all subcubes in $F \setminus G$ are disjoint from $\bigvee G$. Since F is a subcube partition, this means that $\bigcup G = \bigvee G$, which is a subcube. By construction, $1 < |G| < |F|$, and so F is indeed reducible.

To complete the proof, we show that if F is reducible, then the algorithm outputs "reducible". If F is reducible then there is a subset $H \subset F$, with $1 < |H| < |F|$, such that $\bigcup H$ is a subcube. Let $s_i, s_j \in H$, and consider the (i, j) iteration of the outer **for** loop.

We prove inductively that at each iteration of the inner **while** loop, G is contained in H. This holds by construction at the very first step. Now suppose that $G \subseteq H$ and that $s_k \notin G$ intersects $\bigvee G$. Since $G \subseteq H$, also $\bigvee G \subseteq \bigvee H$, and so s_k intersects $\bigvee H$. Since $\bigvee H = \bigcup H$ and the subcubes in F are disjoint, necessarily $s_k \in H$. Hence $G \cup \{s_k\} \subseteq H$.

When the **while** loop ends, all $s_k \notin G$ are disjoint from $\bigvee G$. Since the subcubes in F are disjoint, this means that $\bigvee G = \bigcup G$. Since $G \subseteq H$, necessarily $G \neq F$, and so the algorithm correctly declares that F is reducible. □

2.2 Cubic construction

In this section, we present a construction of an infinite family of tight irreducible subcube partitions.

Theorem 2.8. *Let $n = 2m + 1 \geq 3$. The following subcubes comprise a tight irreducible subcube partition of size $\Theta(n^3)$:*

- *The point* 0^n.
- *All cyclic rotations of* $0^m 1 *^m$.
- *For every* $0 \le i, j, k \le m-1$ *satisfying* $i+j, j+k \le m-1$, *the subcube*

$$0^i 1 *^j 0^k 1 *^{m-1-j-k} 0^j 1 *^{m-1-i-j}.$$

We found this subcube partition by starting with the subcube partition consisting of all rotations of $0^m 1 *^m$ together with all points not covered by them. This subcube partition is reducible, and we can use Algorithm 1 to merge together points into subcubes. One can show inductively that the rotations of $0^m 1 *^m$ never get merged, and so the resulting subcube partition is not trivial. It is precisely the one described in Theorem 2.8.

Here is the resulting partition for $n = 5$:

00000	001∗∗	∗001∗	∗∗001	1∗∗00	01∗∗0
	11∗1∗	1011∗	011∗1	01011	1∗101

Proof of Theorem 2.8. We need to prove three things about the set of subcubes F given in the statement of the theorem: that it is a subcube partition; that it is tight; and that it is irreducible.

Subcube partition The point 0^n covers itself. All other subcubes of F contain at least one 1.

Subcubes of the second type cover *royal points*. These are points which contain a *royal* 1, which is a 1 preceded cyclically by m many 0s. Since $n < 2(m+1)$, there can be at most one royal 1, and so royal points are covered by precisely one subcube of the second type. We will soon see that they are not covered by any subcube of the third type.

We can guarantee that a subcube does not contain any royal point by adding "blocking 1s": if each cyclic interval of length m contains a 1, then the subcube cannot contain any royal point. Each subcube of the third type is contained in the subcube $*^i 1 *^{j+k} 1 *^{m-1-k} 1 *^{m-1-i-j}$, in which the 1s are separated by $j+k, m-1-k, m-1-j \le m-1$ many stars. Consequently, each royal point is covered by precisely one subcube of F.

It remains to handle points $x \ne 0^n$ which are not royal. Let $I+1$ be the index of the first 1 in x. Since x is not royal, $I \le m-1$.

Let $I+1+m+J+1$ be the first 1 in x beyond position $I+1+m$ (so $J \ge 0$). Such a 1 exists since otherwise x starts with $0^I 1$ and ends with 0^{m-I}, and is consequently royal. For the same reason, $J \le m-1$. Since $I+1+m+J+1 \le n$, we see that $I+J \le m-1$.

Let $I+1+J+K+1$ be the first 1 in x beyond position $I+1+J$ (so $K \ge 0$). Such a 1 exists as seen before. Since x is not royal, $J+K \le m-1$. Collecting all the information, we see that x belongs to the subcube of the third type

$$0^I 1 *^J \underset{\substack{\uparrow \\ I+1+J}}{0}{}^K 1 *^{m-1-J-K} \underset{\substack{\uparrow \\ I+1+m}}{0}{}^J 1 *^{m-1-I-J}.$$

If $x \in 0^i 1 *^j 0^k 1 *^{m-1-j-k} 0^j 1 *^{m-1-i-j}$ and we follow the steps above then we find that $i = I$, $j = J$, and $k = K$. Therefore x belongs to a unique subcube of F.

Tightness This is clear, since the subcube 0^n mentions all coordinates.

Irreducibility We need to show that if $G \subset F$ satisfies $1 < |G| < |F|$ then the union of the subcubes in G cannot be a subcube.

Suppose that the union of G is some subcube s. If s', s'' are two distinct subcubes in G, then $s \supseteq s' \vee s''$. It is easy to check that $s' \vee s''$ contains at most two 1s, since the only subcubes in F containing three 1s are those of the third type, and no two of these contain the 1s at the same positions. It follows that s also contains at most two 1s, since it is obtained from $s' \vee s''$ by "erasing" some of the 0/1 coordinates to $*$.

To complete the proof, for each subcube $s \neq *^n$, not contained in any clause of F, and containing at most two 1s, we will exhibit a subcube $t \in F$ which intersects s but is not contained in s. This results in a contradiction: since G intersects t, it must contain t, but then s would contain t. In all cases, t will be a subcube of the second type, that is, a cyclic rotation of $0^m 1 *^m$.

The subcube t intersects the subcube s but is not contained in t if the following two conditions hold:

- $(s_i, t_i) \notin \{(0,1),(1,0)\}$ for all i.
- $(s_i, t_i) \in \{(0,*),(1,*)\}$ for some i.

The coordinate of the second type, or the part containing it, is highlighted in the diagrams below.

Suppose first that s contains no 1s. Since $s \neq 0^n, *^n$, it must contain $*0$ as a cyclic substring. If s ends with $*0$ then we choose $t = *^{m-1}0^m 1*$:

$$\begin{array}{rccc} s = & \dots & * & \boxed{0} \\ t = & *^{m-1}0^m & 1 & \boxed{*} \end{array}$$

If $*0$ is located at a different position, we choose an appropriate subcube of the second type.

Suppose next that s contains a single 1, and furthermore does not contain $1*^m$ as a cyclic substring. If $s_m = 1$ then we choose $t = 0^m 1 *^m$:

$$\begin{array}{rccc} s = & \dots & 1 & \boxed{\neq *^m} \\ t = & 0^m & 1 & \boxed{*^m} \end{array}$$

If the 1 is located at a different position, we choose an appropriate subcube of the second type.

Suppose now that s contains a single 1 and does contain $1*^m$ as a cyclic substring, say it ends with $1*^m$. Since $s \neq 0^m 1 *^m$, necessarily $s_i = *$ for some $i \in [m]$. We choose $t = 0^{i-1}1*^m 0^{m-i+1}$:

$$\begin{array}{rccccc} s = & \dots & * & \dots & \boxed{1} & *^m \\ t = & 0^{i-1} & 1 & *^{m-i} & \boxed{*} & *^{i-1}0^{m-i+1} \end{array}$$

If $1*^m$ is located at a different position, we choose an appropriate subcube of the second type.

Finally, suppose that s contains two 1s, say at positions $i < j$. If $i = m+1$ then we choose $t = 0^m 1 *^m$:

$$\begin{array}{rccccc} s = & \cdots & 1 & \cdots & \boxed{1} & \cdots \\ t = & 0^m & 1 & *^{j-m-2} & \boxed{*} & *^{2m+1-j} \end{array}$$

If the first 1 is located at a different position, we choose an appropriate subcube of the second type (rotate s so that $s_{m+1} = 1$; if $s_i = 1$ for $i \in [m]$, rotate s to the right by $m+1-i$ positions). □

2.3 Minimal size

What is the minimal size of a tight irreducible subcube partition of length n? (The question doesn't make sense without assuming tightness, since $*^n$ is always irreducible.)

When $n = 1$, there is a single tight irreducible subcube partition: $0, 1$. When $n = 2$, there are no tight irreducible subcube partitions. When $n = 3$, there is a unique tight irreducible subcube partition, up to flipping and rearranging coordinates:

$$000, *01, 1*0, 01*, 11.$$

For $n = 4, 5, 6, 7$, Peitl and Szeider [PS22] used a computer search to show that the minimal number of subcubes is $7, 9, 11, 13$, respectively. This is consistent with the following conjecture.

Conjecture 1. *If $n \geq 3$ then the minimal size of a tight irreducible subcube partition of length n is $2n-1$.*

Section 2.3.1 explains the best lower bound on the size, due to Kullmann and Zhao [KZ13]. Sections 2.3.2 and 2.3.3 present two constructions of an infinite family of tight irreducible subcube partitions of length n and size $2n-1$. In Section 2.4 we present several more such constructions which will be useful in Section 3.

2.3.1 Lower bound

Before presenting the constructions of tight irreducible subcube partitions of size $2n-1$, here is the best lower bound on the size, due to Kullmann and Zhao [KZ16]. We give an alternative proof using known results from the literature.

Theorem 2.9. *If $n \geq 4$ then every tight irreducible subcube partition of length n has size at least $n+3$.*

Before proving the theorem, we need a simple lemma.

Definition 2.10 (Regularity)**.** A subcube partition of length n is *regular* if for every $i \in [n]$ and every $b \in \{0,1\}$ there are at least two subcubes $s \in F$ such that $s_i = b$.

This definition is due to Kullmann and Zhao [KZ13], who used the term *nonsingular*. The term *regular* appears in Peitl and Szeider [PS22].

Lemma 2.11. *If F is a tight irreducible subcube partition of length $n \geq 2$ then F is regular.*

Proof. We prove the definition of regularity for $i = 1$.

For $\sigma \in \{0,1,*\}$, let $F_\sigma = \{x : \sigma x \in F\}$. Both $F_0 \cup F_*$ and $F_1 \cup F_*$ are subcube partitions of length $n-1$, and so $\bigcup F_0 = \bigcup F_1$. Since F is tight, F_0, F_1 are non-empty.

If $F_0 = \{x\}$ and $|F_1| > 1$ then the union of the subcubes corresponding to F_1 is the subcube $1x$, contradicting irreducibility.

If $F_0 = \{x\}$ and $|F_1| = 1$ then $F_0 = F_1 = \{x\}$ and so the union of the corresponding subcubes is $*x$. Since F is irreducible, necessarily $x = *^{n-1}$, and so $F = \{0*^{n-1}, 1*^{n-1}\}$. Since F is tight, necessarily $n = 1$, contradicting the assumption $n \geq 2$.

It follows that $|F_0| \geq 2$. Similarly $|F_1| \geq 2$. ◻

We can now prove the size lower bound.

Proof of Theorem 2.9. Let $F = \{s_1, \dots, s_m\}$ be a tight subcube partition of length n. We can identify F with a formula Φ in *conjunctive normal form* (CNF) over variables $x_1, \dots, x_n$ whose clauses are "$x \notin s_i$" for all $i \in [m]$. For example, the subcube partition $0*, 10, 11$ corresponds to the CNF $x_1 \wedge (\bar{x}_2 \vee x_3) \wedge (\bar{x}_2 \vee \bar{x}_3)$.

Since every x belongs to some s_i, the formula Φ is unsatisfiable. It is moreover minimally unsatisfiable, meaning that if we remove any clause, then it becomes satisfiable. Indeed, if we remove the clause "$x \notin s_i$", then any point in s_i would satisfy the formula. Since F is tight, Φ mentions all n variables.

A well-known result attributed to Tarsi [AL86] states that a minimally unsatisfiable CNF mentioning n variables must contain at least $n+1$ clauses, hence $m \geq n+1$.

Suppose that $m = n+1$. Davydov, Davydova, and Kleine Büning [DDKB98, Theorem 12] showed that if a minimally unsatisfiable CNF mentioning n variables contains exactly $n+1$ clauses, then some variable appears once positively and once negatively. In particular, F is not regular, contradicting Lemma 2.11. Hence $m \geq n+2$.

Suppose that $m = n+2$. Kleine Büning [KB00, Theorem 6] showed that there is a unique regular minimally unsatisfiable CNF mentioning n variables which contains exactly $n+2$ clauses, up to renaming and reordering variables. The collection of subcubes corresponding to this CNF consists of $0^n, 1^n$ together with all cyclic rotations of $10*^{n-2}$. When $n \geq 4$, these subcubes are not disjoint: for example, $10*^{n-2}$ and $*^2 10*^{n-4}$ both contain the subcube $1010*^{n-4}$. Hence $m \geq n+3$. ◻

In the following two subsections, we present two constructions of the same sequence of tight irreducible subcube partitions of length $n \geq 3$ and size $2n-1$.

2.3.2 Merging

Our first construction is based on the following lemma, which is used to merge together two subcube partitions.

Definition 2.12 (Reducibility for partial subcube partitions)**.** A subset F' of a subcube partition F of length n is *reducible* if there exists a subset $G \subseteq F'$, with $|G| > 1$, such that the union of the subcubes in G is a subcube different from $*^n$.

Lemma 2.13. *Let F_0, F_1 be two subcube partitions of length n. Let*

$$G = \{0x : x \in F_0 \setminus F_1\} \cup \{1x : x \in F_1 \setminus F_0\} \cup \{*x : x \in F_0 \cap F_1\}.$$

Then

(a) G is a subcube partition of length $n+1$.

(b) If $F_0 \neq F_1$ and at least one of them is tight, then G is tight.

(c) If $F_0 \cap F_1 \neq \emptyset$ and both $F_0 \setminus F_1$ and F_1 are irreducible (or both $F_1 \setminus F_0$ and F_0 are irreducible) then G is irreducible.

Proof. The first two items follow easily from the construction (the condition $F_0 \neq F_1$ in the second item guarantees that the first coordinate is mentioned).

Now suppose that $F_0 \cap F_1 \neq \emptyset$ and both $F_0 \setminus F_1$ and F_1 are irreducible. We need to show that G is irreducible. If not, then there is a subset $H \subset G$, with $1 < |H| < |G|$, whose union is a subcube $x \neq *^n$.

If $x = 0y$ then y is a union of $|H|$ subcubes in $F_0 \setminus F_1$. Since $F_0 \setminus F_1$ is irreducible and $|H| > 1$, necessarily $y = *^n$. However, this contradicts the assumption $F_0 \cap F_1 \neq \emptyset$.

We get a similar contradiction if $x = 1y$, using the irreducibility of F_1.

Finally, if $x = *y$ then y is a union of $|H|$ subcubes of F_0 as well as a union of $|H|$ subcubes of F_1. Since F_1 is irreducible and $y \neq *^n$, necessarily $y \in F_1$. If $y \in F_0$ then $x \in G$, contradicting the assumption $|H| > 1$. If $y \notin F_0$ then $1y \in G$ and so y is a union of subcubes in $F_0 \setminus F_1$. Since $F_0 \setminus F_1$ is irreducible and $y \neq *^n$, necessarily $y \in F_0 \setminus F_1$, contradicting both $y \notin F_0$ and $y \in F_1$. □

We now construct the promised sequence of tight irreducible subcube partitions.

Theorem 2.14. *For each $n \geq 3$ there is a tight irreducible subcube partition S_n of length n and size $2n-1$.*

Proof. We construct the subcube partitions inductively. The starting point is

$$S_3 = \{000, *01, 1*0, 01*, 111\},$$

whose irreducibility can be checked using Algorithm 1. The construction will maintain the invariants that $01*^{n-2} \in S_n$ and $1*^{n-1}, 00*^{n-2} \notin S_n$, and moreover $|S_n| = 2n-1$

Given S_n, we construct S_{n+1} by applying Lemma 2.13 to $F_0 = \{1*^{n-1}, 00*^{n-2}, 01*^{n-2}\}$ and $F_1 = S_n$.

Since F_0 is reducible and F_1 is irreducible, clearly $F_0 \neq F_1$, and so S_{n+1} is tight by Lemma 2.13.

The invariant implies that $F_0 \setminus F_1 = \{1*^{n-1}, 00*^{n-2}\}$ is irreducible. It follows that S_{n+1} is irreducible by Lemma 2.13.

Since $1*^{n-2} \notin F_1$, it follows that $01*^{n-2} \in S_{n+1}$. Since $0*^{n-2} \notin F_0$, it follows that $00*^{n-2} \notin S_{n+1}$. Since $|F_1| > 1$, in particular $*^n \notin F_1$, and so $1*^n \notin S_n$.

Finally, the invariants imply that $F_0 \cap F_1 = \{01*^{n-2}\}$, and so

$$|S_{n+1}| = |F_0 \setminus F_1| + |F_1 \setminus F_0| + |F_0 \cap F_1| = 2 + (|F_1| - 1) + 1 = 2n + 1,$$

using $|F_1| = 2n - 1$. □

Here are the resulting subcube partitions for $n = 3, 4, 5$:

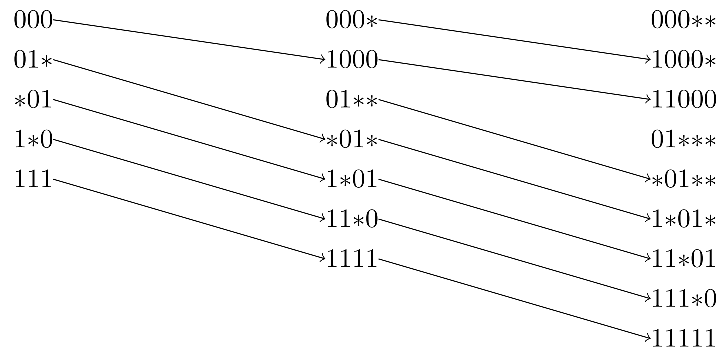

2.3.3 Twisting

Our second construction starts with the observation

$$x00 \cup x1{*} = x{*}0 \cup x11.$$

Up to permutation and flipping of coordinates, this is the only way in which a set of points can be written as a union of two subcubes in two different ways (we leave the proof to the reader). Following Kullmann and Zhao [KZ16, Definitions 45–46], we call such a pair of subcubes an *nfs-pair*. The *nfs-flip* of the pair on the left is the pair on the right.

Definition 2.15 (Nfs-pair, nfs-flip)**.** Two subcubes s, t constitute an *nfs-pair* if they differ on exactly two positions i, j, where $(s_i, t_i) \in \{(0, 1), (1, 0)\}$ and $t_j = *$.

The *nfs-flip* of s, t is the pair of subcubes s', t' obtained by copying the coordinates except for i, j, and setting $s'_i = *$, $s'_j = s_j$, $t'_i = t_i$, $t'_j = \bar{s}'_j$.

Lemma 2.16. *If s, t is an nfs-pair with nfs-flip s', t' then $s \cup t = s' \cup t'$.*

The construction is based on the following simple corollary of Lemma 2.13.

Lemma 2.17. *Let F be a tight irreducible subcube partition containing an nfs-pair s, t, and let s', t' be its nfs-flip. The following subcube partition is tight and irreducible, for any $b \in \{0, 1\}$:*

$$G = \{{*}x : x \in F, x \neq s, t\} \cup \{bs, bt, \bar{b}s', \bar{b}t'\}.$$

Furthermore, $|G| = |F| + 2$, and G contains the nfs-pairs bs, bt and $\bar{b}s', \bar{b}t'$.

Proof. Let F' be the formula obtained from F by replacing s, t with s', t'. We apply Lemma 2.13 on $F_b = F$ and $F_{\bar{b}} = F'$, obtaining the stated subcube partition G.

Since $F \neq F'$ and F is tight, G is tight.

Clearly F cannot consist only of s, t, and so $F \cap F' \neq \emptyset$. Since F and $F' \setminus F \subset F'$ are both irreducible, it follows that G is irreducible. ☐

In order to obtain the sequence S_n constructed in Theorem 2.14 using Lemma 2.17, start with S_3. Given S_n, apply the lemma with $s = 1^n$, $t = 1^{n-2}{*}0$, and $b = 1$, and rotate the resulting subcube partition once to the left. The result is S_{n+1}. Here is an example:

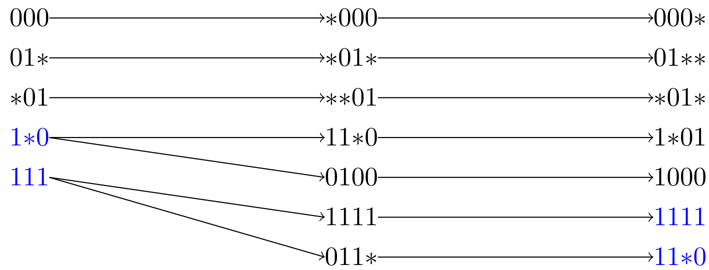

2.4 Minimal weight

In Section 3, we will consider irreducible subcube partitions over larger alphabets. As we show in Section 3.1, one of the ways to construct an irreducible subcube partition over an alphabet $\{0, \dots, q-1\}$ is to start with an irreducible subcube partition over $\{0, 1\}$, and replace each 1 in each subcube with each of $\{1, \dots, q-1\}$. The resulting number of subcubes is

$$\sum_{s \in F} (q-1)^{\#_1(s)},$$

where F is the subcube partition we start with, and $\#_1(s)$ is the number of 1s in s. This suggests looking for a tight irreducible subcube partition which minimizes the above objective function.

The concept of *majorization* allows us to optimize this objective function for all q's at once.

Definition 2.18 (Weight vector)**.** Let F be a subcube partition of length n. Its *weight vector* is the vector $w(F) = w_0, \dots, w_n$, where w_h is the number of subcubes of F of *weight* h, that is, with h many 1s.

The notation $w_{\geq h}$ stands for $w_h + \cdots + w_n$, which is the number of subcubes with at least h many 1s.

Definition 2.19 (Majorization)**.** Let a, b be two weight vectors of length $n+1$. We say that *a majorizes b* if for every $h \leq n$, we have $a_{\geq h} \geq b_{\geq h}$.

Lemma 2.20. *Let F, G be subcube partitions of length n. If $w(F)$ majorizes $w(G)$ then for all $q \geq 2$,*

$$\sum_{s \in F} (q-1)^{\#_1(s)} \geq \sum_{s \in G} (q-1)^{\#_1(s)}.$$

Proof. Let $r = q - 1 \geq 1$. We will show that

$$\sum_{h=0}^{n} w_h(F) r^h \geq \sum_{h=0}^{n} w_h(G) r^h.$$

Indeed,

$$\sum_{h=0}^{n} w_h(F) r^h = w_{\geq 0}(F) + \sum_{h=1}^{n} w_{\geq h}(F)(r^h - r^{h-1}) \geq w_{\geq 0}(G) + \sum_{h=1}^{n} w_{\geq h}(G)(r^h - r^{h-1}) = \sum_{h=0}^{n} w_h(G) r^h. \quad \square$$

Lemma 2.20 allows us to reformulate our goal: find the minimal weight vectors (in the sense of majorization) of the tight irreducible subcube partitions of length n. (There could be more than one minimal weight vector since majorization is not a linear order.)

Conjecture 2. *For every $n \geq 3$, the minimal weight vectors of tight irreducible subcube partitions of length n are $1, n-1, n-1, 0, \dots, 0$ and $1, n, n-3, 1, 0, \dots, 0$.*

In Section 2.4.1, we show that Conjecture 1 implies the lower bound part of Conjecture 2. In Section 2.4.2 we give matching constructions.

Unconditionally, we can show that every tight irreducible subcube partition of length $n \geq 3$ must contain a subcube of weight 2.

Lemma 2.21. *If F is a tight irreducible subcube partition of length $n \geq 3$ then F contains a subcube of weight at least 2.*

Proof. Suppose that every subcube in F has weight at most 1. Let $s \in F$ be the subcube containing 1^n. If s has weight 0 then $s = *^n$, contradicting the tightness of F. If s has weight 1 then, without loss of generality, $s = 1*^{n-1}$. The union of all other subcubes of F must be $0*^{n-1}$, and so by irreducibility, $F = \{0*^{n-1}, 1*^{n-1}\}$, contradicting tightness. □

2.4.1 Lower bound

In this section we prove the lower bound part of Conjecture 2, assuming Conjecture 1. As we explain in the proof, this amounts to ruling out the weight vector $1, n, n-2, 0, \ldots, 0$.

Theorem 2.22. *Assume Conjecture 1. For every $n \geq 3$, the weight vector of any tight irreducible subcube partition of length n majorizes either $1, n-1, n-1, 0, \ldots, 0$ or $1, n, n-3, 1, 0, \ldots, 0$.*

Proof. Let F be a tight irreducible subcube partition of length n, and let w be its weight vector. The theorem states that (i) $w_{\geq 0} \geq 2n-1$; (ii) $w_{\geq 1} \geq 2n-2$; and either (iii) $w_{\geq 2} \geq n-1$ or (iv) $w_{\geq 2} \geq n-2$ and $w_{\geq 3} \geq 1$.

We start with the following observation: $w_h \leq \binom{n}{h}$. Indeed, every subcube s of weight h contains the point x_s obtained by switching all $*$s to 0s, which has weight h. Since the subcubes in F are disjoint, every $s \in F$ of weight h has a different x_s. Since there are $\binom{n}{h}$ many possible x_s, it follows that $w_h \leq \binom{n}{h}$.

The inequality $w_{\geq 0} \geq 2n-1$ is Conjecture 1. Since $w_0 \leq 1$, the inequality $w_{\geq 1} \geq 2n-2$ follows. Since $w_1 \leq n$, we deduce the inequality $w_{\geq 2} \geq n-2$. To complete the proof, we need to show that either (iii) $w_{\geq 2} \geq n-1$ or (iv) $w_{\geq 3} \geq 1$. We will show that the assumptions $w_{\geq 2} = n-2$ and $w_{\geq 3} = 0$ lead to a contradiction.

Suppose, therefore, that $w_2 = n-2$ and $w_{\geq 3} = 0$. Since $w_{\geq 0} \geq 2n-1$ and $w_0 \leq 1$, $w_1 \leq n$, this implies that $w_0 = 1$ and $w_1 = n$.

Since $w_1 = n$, for every $i \in [n]$ there is a subcube $s^{(i)} \in F$ which contains 1 in the i'th position: $s^{(i)}_i = 1$. The point 0^n is covered by the unique subcube $s^{(0)} \in F$ of weight 0. Since $s^{(0)}$ and $s^{(i)}$ must conflict, necessarily $s^{(0)}_i = 0$ (this is the only possible conflict), and so $s^{(0)} = 0^n$ is a point.

Since $0^n \in F$, the subcubes $s^{(i)}$ cannot be points. Indeed, if $s^{(i)}$ is a point then $s^{(i)} = 0^{i-1}10^{n-i}$, and so $s^{(0)} \cup s^{(i)} = 0^{i-1}*0^{n-i}$, contradicting irreducibility. Consequently, all points in F have even weight.

Every subcube in F which is not a point contains an equal number of points of even weight and of odd weight. In contrast, all points in F have even weight. Since $\{0,1\}^n$ contains an equal number of points of either parity, we reach a contradiction. □

2.4.2 Construction

In this section, we prove (unconditionally) the upper bound part of Conjecture 2, by constructing tight irreducible subcube partitions of length $n \geq 3$ and weight vectors $1, n-1, n-1, 0, \ldots, 0$ and $1, n, n-3, 1, 0, \ldots, 0$. The constructions will use the methods of Section 2.3.2. The same subcube partitions can also be constructed using the methods Section 2.3.3; we leave the details to the reader.

Theorem 2.23. *For each $n \geq 3$ there is a tight irreducible subcube partition A_n whose weight vector is $1, n-1, n-1, 0, \ldots, 0$.*

Proof. We construct the subcube partitions inductively, starting with

$$A_3 = \{*00, 001, 01*, 110, 1*1\},$$

which is obtained from S_3 of Theorem 2.14 by flipping the third coordinate. The construction will maintain the invariants that $01*^{n-2} \in A_n$ and $1*^{n-1}, 00*^{n-2} \notin A_n$.

Given A_n, we construct A_{n+1} by applying Lemma 2.13 to $F_0 = A_n$ and $F_1 = \{1*^{n-1}, 00*^{n-2}, 01*^{n-2}\}$, and rotating the result G once to the left, that is, $A_{n+1} = \{xb : bx \in G\}$, where $b \in \{0, 1, *\}$ and $x \in \{0, 1, *\}^n$.

Since F_0 is irreducible and F_1 is reducible, clearly $F_0 \neq F_1$, and so A_{n+1} is tight by Lemma 2.13.

The invariant implies that $F_1 \setminus F_0 = \{1*^{n-1}, 00*^{n-2}\}$ is irreducible, and so A_{n+1} is irreducible by Lemma 2.13.

The invariant states that $01*^{n-2} \in F_0$. Since $01*^{n-2} \in F_1$ by definition, it follows that $*01*^{n-2} \in G$, and so $01*^{n-1} \in A_{n+1}$. Since A_{n+1} is irreducible, necessarily $00*^{n-1} \notin A_{n+1}$ (otherwise $00*^{n-1} \cup 01*^{n-1} = 0*^n$ would be a subcube) and $1*^n \notin A_{n+1}$ (otherwise the union of all other subcubes would be $0*^n$).

Finally, the invariants imply that $F_1 \setminus F_0 = \{1*^{n-1}, 00*^{n-2}\}$, and so compared to F_0, the subcube partition G gains one subcube of weight 2 (namely, $11*^{n-1}$) and one subcube of weight 1 (namely, $100*^{n-2}$); all other subcubes originate from F_0 and maintain their weight. □

Theorem 2.24. *For each $n \geq 3$ there is a tight irreducible subcube partition D_n whose weight vector is $1, n, n-3, 1, 0, \ldots, 0$.*

Proof. The proof is very similar to that of Theorem 2.23. We take

$$D_3 = S_3 = \{*01, 000, 01*, 111, 1*0\}.$$

The rest of the proof is identical. □

The subcube partition D_3 is obtained from A_3 by flipping the third coordinate, and this holds for every n, by construction. By flipping coordinates appropriately, we can also obtain other tight irreducible subcube partitions B_n, C_n whose weight vectors are $1, n-1, n-1, 0, \ldots, 0$. The subcube partition B_n is obtained by flipping the first coordinate, and C_n is obtained by flipping both the first and the third coordinates. The subcube partitions B_n, C_n can also be obtained using an iterative construction as above.

Here are the subcube partitions A_5, B_5, C_5, D_5 after rotation once to the left:

0000∗	0000∗	0100∗	0100∗
01000	01001	00001	00000
0∗100	0∗101	0∗101	0∗100
0∗∗10	0∗∗11	0∗∗11	0∗∗10
1∗∗∗0	1∗∗∗1	1∗∗∗1	1∗∗∗0
10001	10000	11000	11001
∗1001	∗1000	∗0000	∗0001
∗∗101	∗∗100	∗∗100	∗∗101
∗∗∗11	∗∗∗10	∗∗∗10	∗∗∗11
A_5	B_5	C_5	D_5

Among all subcube partitions obtained from A_n by flipping coordinates, these are the only ones whose weight vector is either $1, n-1, n-1, 0, \ldots, 0$ or $1, n, n-3, 1, 0, \ldots, 0$.

2.5 Maximal number of points

In the following section, we tackle the problem of maximizing the number of subcubes in an irreducible subcube partition. As a warm-up, we start with the problem of maximizing the number of points (zero-dimensional subcubes) in an irreducible subcube partition.

For $n = 3, 4, 5, 6$, a computer search reveals that the maximum number of points in an irreducible subcube partition of length n is $2, 4, 8, 16$. This is consistent with the following conjecture.

Conjecture 3. *If $n \geq 3$ then the maximum number of points in an irreducible subcube partition of length n is 2^{n-2}.*

It is easy to see that an irreducible subcube partition of length n contains at most 2^{n-1} points. Indeed, if the subcube partition contained more than 2^{n-1} points then there would be two points differing in a single coordinate. The union of these two points is an edge (a one-dimensional subcube), contradicting irreducubility.

In the rest of this section, we construct an irreducible subcube partition of length n containing 2^{n-2} many points. The construction uses the following lemma.

Lemma 2.25. *If $F \neq \{0*^{n-1}, 1*^{n-1}\}$ is an irreducible subcube partition of length n then the following is an irreducible subcube partition of length $n+1$:*

$$G = \{00t, 11t : 0t \in F\} \cup \{01t, 10t : 1t \in F\} \cup \{**t : *t \in F\}.$$

Proof. Let F' be the irreducible subcube partition obtained from F by flipping the first coordinate. The subcube partition G results from applying Lemma 2.13 to F and F'. According to the lemma, in order to show that G is irreducible, it suffices to show that $F \cap F' \neq \emptyset$.

If $F \cap F' = \emptyset$ then all subcubes in F start with 0 or 1. Hence the union of all subcubes in F starting with 0 is $0*^{n-1}$, and the union of all subcubes in F starting with 1 is $1*^{n-1}$. Since F is irreducible, it follows that $F = \{0*^{n-1}, 1*^{n-1}\}$, contrary to the assumption. Therefore $F \cap F' \neq \emptyset$, completing the proof. □

Corollary 2.26. *If $F \neq \{0*^{n-1}, 1*^{n-1}\}$ is an irreducible subcube partition of length n then for every $k \geq 1$, the following is an irreducible subcube partition of length $n+k$:*

$$G = \{a_1 \ldots a_k t : bt \in F, a_1, \ldots, a_k, b \in \{0,1\}, a_1 \oplus \cdots \oplus a_k = b\} \cup \{*^k t : *t \in F\}.$$

Proof. Apply the lemma iteratively $k-1$ times to F, noticing that the subcube partition constructed in the lemma is never of the form $\{0*^m, 1*^m\}$. □

In order to construct an irreducible subcube partition of length n with 2^{n-2} many points, we apply the corollary to the irreducible subcube partition S_3 from Theorem 2.14.

Theorem 2.27. *For every $n \geq 3$ there is an irreducible subcube partition of length n containing 2^{n-2} many points.*

Proof. The subcube partition $S_3 = \{000, *01, 1*0, 01*, 111\}$ is irreducible (according to Theorem 2.14) and contains two points. Applying Corollary 2.26 with $k = n-3$, we get an irreducible subcube partition of length n in which each of the two original points gives rise to 2^{n-3} points, for a total of 2^{n-2} points. □

Using similar ideas, we can construct an irreducible subcube partition of length n containing any even number of points between 2 and 2^{n-2}. We leave the details to the reader.

2.6 Maximal size

What is the maximal size of an irreducible subcube partition of length n? Here are some experimental results:

n	3	4	5	6	7	8	9
Lower bound	5	9	20	40	80	160	320
Upper bound	5	9	20	40	83	166	334

The lower bounds correspond to explicit constructions. We describe such constructions for $n = 3,4,5,6,7$ in Section 2.6.2. The upper bounds are a combination of experimental results and an argument appearing in Section 2.6.1, which gives an upper bound of $\frac{2n-1}{3n-1}2^n \approx \frac{16}{3}2^{n-3}$.

We make the following conjecture, which is consistent with the table.

Conjecture 4. *For every $n \geq 5$, the maximal size of an irreducible subcube partition of length n is $5 \cdot 2^{n-3}$.*

The value $5 \cdot 2^{n-3}$ is best possible, assuming Conjecture 3.

Lemma 2.28. *Assume Conjecture 3. For every $n \geq 3$, every irreducible subcube partition of length n has size at most $5 \cdot 2^{n-3}$.*

Proof. Let F be an irreducible subcube partition of length n. According to Conjecture 3, F contains $m \leq 2^{n-2}$ points. All other subcubes of F cover at least two points, and so the size of F is at most

$$m + \frac{2^n - m}{2} = 2^{n-1} + \frac{m}{2} \leq 2^{n-1} + 2^{n-3}. \qquad \square$$

2.6.1 Upper bound

In this section, we use a result of Forcade [For73] to give an upper bound on the size of irreducible subcube partitions.

Theorem 2.29. *For every $n \geq 3$, the size of any irreducible subcube partition of length n is at most*

$$\frac{2n-1}{3n-1} 2^n = \left(\frac{16}{3} - \Theta\left(\frac{1}{n}\right)\right) 2^{n-3}.$$

Proof. Let F be an irreducible subcube partition of length n. Let G be a subcube partition obtained from F by subdividing each subcube of dimension larger than 1 into edges (subcubes of dimension 1) in an arbitrary way.

Since F is irreducible, no two points in G span an edge. Therefore the set of edges in G constitutes a maximal matching in the n-dimensional hypercube. Forcade [For73] proved that any maximal matching in the n-dimensional hypercube contains $m \geq \frac{n}{3n-1} 2^n$ edges. Therefore

$$|F| \leq |G| = m + (2^n - 2m) = 2^n - m \leq 2^n - \frac{n}{3n-1} 2^n. \qquad \square$$

Forcade showed that the bound $\frac{n}{3n-1} 2^n$ is asymptotically tight by giving a matching construction. Therefore this method cannot prove the conjectured upper bound $5 \cdot 2^{n-3}$.

2.6.2 Constructions

In this section we give irreducible subcube partitions of maximal size for lengths $3, 4, 5, 6, 7$. In each case, irreducibility can be proved using Algorithm 1. We chose to present these constructions since we can describe them compactly.

When $n = 3$, we can take S_3 of Theorem 2.14.

For $n = 4$, we have the following construction:

$$aabc\colon a = b \oplus c$$
$$*1aa$$
$$0*a\bar{a}$$
$$10**$$

Here $a, b, c \in \{0, 1\}$; we use similar notation below. Written out in full, this is

$$0000, 0011, 1101, 1110, *100, *111, 0*01, 0*10, 10**.$$

For $n = 5$, we have the following construction:

$$\begin{aligned}&abcde\colon a \oplus b = c = d \oplus e\\ &{*}a0\bar{a}a\\ &\bar{a}{*}1aa\\ &ab{*}a\bar{b}\\ &aa1{*}a\\ &\bar{a}a0a{*}\end{aligned}$$

For $n = 6$, we have the following construction:

$$\begin{aligned}&abcdef\colon a \oplus b = c \oplus d = e \oplus f\\ &{*}ba\bar{b}ab\\ &a{*}\bar{a}bab\\ &ab{*}b\bar{a}b\\ &aba{*}a\bar{b}\\ &ab\bar{a}b{*}\bar{b}\\ &aba\bar{b}\bar{a}{*}\end{aligned}$$

Here are two constructions for $n = 7$:

$$\begin{aligned}&abcdefg\colon f = b \oplus c \oplus d \oplus e, g = a \oplus b \oplus c \oplus d\\ &0a{*}bb\bar{c}a\colon c = a \oplus b\\ &ab00a{*}\bar{c}\colon c = a \oplus b\\ &ab00\bar{a}c{*}\colon c = a \oplus b\\ &ab0{*}1\bar{c}\bar{b}\colon c = a \oplus b\\ &{*}ab11\bar{c}c\colon c = a \oplus b\\ &{*}ab\bar{b}0a\bar{a}\\ &aba\bar{a}a{*}c\colon c = a \oplus b\\ &ab1a{*}bb\\ &ab1a\bar{a}\bar{b}{*}\\ &a{*}1\bar{a}0b\bar{c}\colon c = a \oplus b\\ &ab\bar{a}1{*}\bar{b}\bar{b}\\ &1ab{*}bca\colon c = a \oplus b\end{aligned}$$

$$\begin{aligned}&abcdefg\colon f = b \oplus c \oplus d \oplus e, g = a \oplus b \oplus c \oplus d\\ &0a{*}b\bar{c}\bar{b}\bar{a}\colon c = a \oplus b\\ &0abb{*}b\bar{a}\\ &1ab{*}\bar{c}b\bar{a}\colon c = a \oplus b\\ &1abbc{*}a\colon c = a \oplus b\\ &{*}abd\bar{e}ba\colon e = a \oplus d\\ &abdeg\bar{d}{*}\colon e = a \oplus d, g = a \oplus b \oplus d\\ &abd\bar{d}\bar{g}{*}c\colon c = a \oplus b, g = a \oplus b \oplus d\\ &a{*}b\bar{c}d\bar{b}g\colon c = a \oplus b, g = a \oplus b \oplus d\end{aligned}$$

2.7 Maximal minimum dimension

All irreducible subcube partitions we have exhibited so far contain points. Is this necessary? More generally, given n, what is the maximal d such that there exists a tight irreducible subcube partition in which every subcube has dimension at least d? (The question doesn't make sense without assuming tightness, since $*^n$ is always irreducible.)

The constructions we give below suggest the following conjecture.

Definition 2.30 (Minimum dimension)**.** For a subcube partition F, let $\delta(F)$ denote the minimum dimension of a subcube of F.

Conjecture 5. *Every tight irreducible subcube partition F of length n satisfies $\delta(F) \leq n/2 - o(n)$.*

One can similarly ask for the maximum value of $\Delta(F)$, which is the minimum codimension of a subcube of F, over all irreducible formulas of length n. If F has length n then clearly $\Delta(F) \leq n-1$ (for $n \geq 3$). Conversely, Section 2.6 suggests that there are irreducible formulas of length n with $\Delta(F) = n-1$.

In the remainder of this section, we give a construction matching Conjecture 5. The construction is based on the following lemma.

Lemma 2.31. *Let F be a subcube partition of length n. Define*

$$G = \{*t** : t* \in F\} \cup \{0tb*, 1t*b : b \in \{0,1\}, tb \in F\}.$$

Then

(a) G is a subcube partition of length $n+2$.

(b) If F is tight then G is tight.

(c) If F is irreducible and contains a subcube ending with a star then G is irreducible and contains a subcube ending with a star.

(d) $\delta(G) \geq \delta(F)+1$. Furthermore, if $\delta(F)$ is attained at a subcube not ending with a star then the same holds for G, and moreover $\delta(G) = \delta(F)+1$.

Proof. Let $F_0 = \{tc* : tc \in F\}$ and $F_1 = \{t*c : tc \in F\}$, where $c \in \{0,1,*\}$ in both cases. Applying Lemma 2.13 to F_0, F_1, we obtain the subcube partition G.

Suppose that F is tight. For every $i \in \{1,\ldots,n-1\}$, some subcube of F mentions coordinate i. The corresponding subcube or subcubes of G mention coordinate $i+1$. Some subcube of F mentions coordinate n. The corresponding subcubes of G mention the remaining coordinates $1, n+1, n+2$.

Suppose that F is irreducible and contains a subcube $s \in F$ ending with a star. The irreducibility of F directly implies the irreducibility of F_0 and F_1. Since $s* \in F_0 \cap F_1$, Lemma 2.13 implies that G is irreducible. Furthermore, $*s* \in G$ is a subcube ending with a star.

The remaining claims are easy to verify directly once we notice that the dimension of a subcube is the number of star coordinates. □

We apply the construction on three specific tight irreducible subcube partitions (one only for $n = 4$) in order to obtain the following result, which gives the best constructions we are aware of.

Theorem 2.32. *For every odd $n \geq 3$ there is a tight irreducible subcube partition F of length n with $\delta(F) = \frac{n-3}{2}$.*

For $n = 4$ there is a tight irreducible subcube partition F of length n with $\delta(F) = \frac{n-4}{2}$.

For every even $n \geq 6$ there is a tight irreducible subcube partition F of length n with $\delta(F) = \frac{n-2}{2}$.

Proof. The first part follows from applying Lemma 2.31 to the tight irreducible subcube partition S_3 of Theorem 2.14. The second part follows from taking the tight irreducible subcube partition S_4 of the same theorem. The third part follows from applying Lemma 2.31 to the following tight irreducible subcube partition, whose irreducibility can be checked using Algorithm 1:

0*0*1*	00**0*	001*1*	010*0*	0110**
1**0*1	10***0	10*1*1	11*0*0	1101**
				*111**

□

2.8 Homogeneous subcube partitions

So far we have considered various parameters of irreducible subcube partitions, attempting to optimize them. The final question we consider concerns subcube partitions in which all subcubes have the same codimension.

Definition 2.33 (Homogeneity)**.** An *(n,k)-homogeneous* subcube partition is a tight subcube partition of length n in which all subcubes have codimension k.

In this section, we explore the following question: for which n, k does there exist an irreducible (n,k)-homogeneous subcube partition?

Here is a table with some experimental results:

	$n=4$	$n=5$	$n=6$	$n=7$	$n=8$	$n=9$
$k=3$	Y	N	N	N	N	N
$k=4$		N	Y	N	N	N
$k=5$			Y	Y	Y	Y
$k=6$				Y	Y	Y
$k=7$					Y	Y
$k=8$						Y

This leads to the following conjecture.

Conjecture 6. *For every $k \geq 5$ there exists an irreducible $(k+1,k)$-homogeneous subcube partition. (In this subcube partition, all subcubes are edges.)*

In Section 2.8.1 we prove several elementary results: an irreducible $(n,1)$-homogeneous subcube partition exists only for $n=1$; no irreducible $(n,2)$-homogeneous partition exists; and for $k \geq 3$, if an irreducible (n,k)-homogeneous subcube partition exists then $k+1 \leq n \leq 2^k - 3$.

In Section 2.8.2 we show that the weight distribution of an (n,k)-homogeneous subcube partition is binomial.

In Section 2.8.3, we present a construction of two infinite families of irreducible homogeneous subcube partitions.

Finally, in Section 2.8.4 we show that an irreducible $(n,3)$-homogeneous partition exists only for $n=4$, and in Section 2.8.5 we show that an irreducible $(n,4)$-homogeneous partition exists only for $n=6$.

2.8.1 Elementary bounds

We start with the following general bound.

Lemma 2.34. *Suppose that $n \geq 4$ and $k \geq 2$. If there exists an irreducible (n,k)-homogeneous subcube partition then $k+1 \leq n \leq 2^k - 3$.*

Proof. Let F be an irreducible (n,k)-homogeneous subcube partition. Clearly $n \geq k$. If $n = k$ then all subcubes in F are points, and so F is not irreducible. Hence $n \geq k+1$. Since F has size 2^k, the upper bound $n \leq 2^k - 3$ follows from Theorem 2.9. □

We now determine when an irreducible (n,k)-homogeneous subcube partition exists for $k=1$ and $k=2$.

Lemma 2.35. *If F is an irreducible $(n,1)$-homogeneous subcube partition then $n=1$ and $F = \{0,1\}$.*

Proof. The two subcubes in F contain a single non-star position, which must be identical. Since F is tight, necessarily $n=1$, and so $F=\{0,1\}$. □

Lemma 2.36. *There are no irreducible $(n,2)$-homogeneous subcube partitions, for any n.*

Proof. If $n \geq 4$ then Lemma 2.34 implies that $3 \leq n \leq 1$, which is impossible. If $n=2$ then all subcubes in the subcube partition are points, and so it cannot be irreducible. It remains to rule out the case $n=3$.

Suppose that F is an irreducible $(3,2)$-homogeneous subcube partition. Without loss of generality, $00{*} \in F$. The subcube covering 010 must be ${*}10$, since $01{*}$ would violate irreducibility and $0{*}0$ intersects $00{*}$. Similarly, the subcube covering 100 must be $1{*}0$. However, the subcubes ${*}10$ and $1{*}0$ intersect. □

2.8.2 Weight distribution

In this section we prove the following surprising property, which involves the concept of weight vector defined in Section 2.4.

Lemma 2.37. *The weight vector of any* (n, k)*-homogeneous subcube partition is*

$$\binom{k}{0}, \binom{k}{1}, \ldots, \binom{k}{k}, 0, \ldots, 0.$$

Proof. Let F be an (n, k)-homogeneous subcube partition, and let w be its weight vector. Considering the number of points of weight ℓ which are covered, for each $\ell \in \{0, \ldots, k\}$ we have

$$\sum_{r=0}^{\ell} \binom{n-k}{\ell - r} w_r = \binom{n}{\ell}.$$

This is a triangular system of equations, and so it has a unique solution. In other words, all (n, k)-homogeneous subcube partitions (if any) have the same weight vector.

The argument above applies even if we don't assume that F is tight. Therefore all (n, k)-homogeneous subcube partitions have the same weight vector as the subcube partition $\{x{*}^{n-k} : x \in \{0, 1\}^k\}$, whose weight vector is the one in the statement of the lemma. □

2.8.3 Infinite families

In this section we construct two infinite families of irreducible homogeneous subcube partitions, using the following lemma.

Lemma 2.38. *Let* $k \geq 2$*. If there exists an irreducible* (n, k)*-homogeneous subcube partition then there exists an irreducible* $(3n, 2k)$*-homogeneous subcube partition.*

Proof. Let F be an irreducible (n, k)-homogeneous subcube partition. We start by observing that for each coordinate i, there must be some subcube $s \in F$ with $s_i = *$. Otherwise, the union of the subcubes $s \in F$ with $s_i = 0$ will be ${*}^{i-1}0{*}^{n-i}$, which contradicts irreducibility.

Repeat the following operation n times to F: apply Lemma 2.31, and rotate the result twice to the right (equivalently, replace $t* \in F$ with $***t$ and $tb \in F$ with $b{*}0t, {*}b1t$, where $b \neq *$). The observation in the preceding paragraph ensures that the resulting subcube partition G is tight and irreducible. By construction, G has length $n + 2n = 3n$ and codimension $k + k = 2k$. □

Using appropriate starting points, we obtain two infinite families of irreducible homogeneous subcube partitions.

Theorem 2.39. *For every* $t \geq 0$ *there are irreducible* $(3^t \cdot 4, 2^t \cdot 3)$*-homogeneous subcube partitions as well as irreducible* $(3^t \cdot 6, 2^t \cdot 4)$*-homogeneous subcube partitions.*

Proof. Apply Lemma 2.38 to the following two homogeneous subcube partitions, whose irreducibility can be verified using Algorithm 1:

$$\{000*, 1{*}00, 01{*}0, {*}010, 10{*}1, {*}101, 0{*}11, 111*\},$$
$$\{0000{*}{*}, 01{*}{*}00, 001{*}{*}1, 01{*}01{*}, 0{*}01{*}1, 0{*}{*}110, 10{*}11{*}, 10{*}{*}01,$$
$$1101{*}{*}, 111{*}{*}0, 1{*}00{*}0, 1{*}{*}011, {*}010{*}0, {*}0{*}100, {*}111{*}1, {*}1{*}001\}.$$
□

2.8.4 Codimension 3

Theorem 2.39 shows that an irreducible $(4,3)$-homogeneous subcube partition exists. In this section, we show that an irreducible $(n,3)$-homogeneous subcube partition exists only for $n = 4$, and that it is unique up to permutation and flipping of coordinates.

Theorem 2.40. *If there exists an irreducible* $(n,3)$*-homogeneous subcube partition then* $n = 4$*.*

Proof. Let F be an irreducible $(n,3)$-homogeneous subcube partition. If $n = 3$ then all subcubes in F are points, and so F is irreducible. Therefore $n \geq 4$.

Every two subcubes in F must conflict, and so they must share at least one non-star position. We show that they must in fact share at least two non-star positions.

Indeed, assume to the contrary that $s,t \in F$ share exactly one non-star position, without loss of generality $s = 000{*}{*}{*}^{n-5}$ and $t = {*}{*}100{*}^{n-5}$. Let $u \in F$ be the subcube covering 100^{n-2}, and let $v \in F$ be the subcube covering 010^{n-2}. By construction $u, v \neq s,t$. Also, $u \neq v$, since otherwise $u \supseteq 100^{n-2} \vee 010^{n-2} = {*}{*}0^{n-2}$, and so u intersects s.

Since s conflicts with u,v, we have $u_1 = 1$ and $v_2 = 1$. Since t conflicts with u,v, we have $u_3 = v_3 = 0$. If $u_2 = 0$ then $u = 100{*}^{n-3}$ and so $s \cup u = {*}00{*}^{n-3}$ is a subcube, contradicting irreducibility. Therefore $u_1u_2u_3 = 1{*}0$, and the remaining non-star in u is 0. Similarly, $v_1v_2v_3 = {*}10$, and the remaining non-star in v is 0. But then u cannot conflict with v, and we have reached a contradiction.

The argument above shows that any two subcubes in F share at least two non-star positions. Without loss of generality, $s = 000{*}^{n-3} \in F$. Let $t \in F$ be the subcube containing the point 10^{n-1}. Irreducibility implies that $t \neq 100{*}^{n-3}$ and disjointness implies that t doesn't start with ${*}00$, and so $t = 10{*}0{*}^{n-4}$ without loss of generality.

Lemma 2.11 shows that F is regular (see Definition 2.10). If $n \geq 5$ then there exist two subcubes $u,v \in F$ with $u_5 = v_5 = 0$. Since u,v each share two non-star positions with both s and t, these must be the first two positions. Since u,v must conflict with s,t, their first two positions must be either 01 or 11, and so $u,v \in \{01{*}{*}0{*}^{n-5}, 11{*}{*}0^{n-5}\}$. However, $u \cup v = {*}1{*}{*}0{*}^{n-5}$ is a subcube, contradicting irreducibility. We conclude that $n = 4$. □

We prove the uniqueness of the subcube partition in Theorem 2.39 using Lemma 2.37.

Theorem 2.41. *There is a unique irreducible* $(4,3)$*-homogeneous subcube partition, up to permutation and flipping of coordinates.*

Proof. Let F be an irreducible $(4,3)$-homogeneous subcube partition. According to Lemma 2.37 there is a unique subcube of weight 0 in F, say $000{*} \in F$.

The lemma also shows that F has three subcubes of weight 1. Since they are disjoint from one another, each of them must contain a 1 in a different position. Since they are disjoint from $000{*}$, the 1 cannot appear in position 4, and so the 1s must appear in positions $1, 2, 3$.

The unique star in each of the subcubes of weight 1 cannot appear in the last position, since otherwise the union with $000{*}$ would also be a subcube. Therefore without loss of generality, $1{*}00 \in F$. The subcube with 1 in the second position cannot be ${*}100$ since this intersects $1{*}00$, and so $01{*}0 \in F$. Repeating the argument, ${*}010 \in F$.

Lemma 2.37 shows that F has three subcubes of weight 2. Since they are disjoint from $1{*}00, 01{*}0, {*}010$, the last position must contain 1 (for example, $11{*}0$ intersects $1{*}00$). The unique 0 must be in different positions, and so the 0s must appear in positions $1, 2, 3$. The subcube of the form $0??1$ cannot be $01{*}1$ since otherwise $01{*}0 \cup 01{*}1 = 01{*}{*}$ would contradict irreducibility, and so $0{*}11 \in F$. Similarly $10{*}1, {*}101 \in F$.

Finally, the remaining points uncovered by the subcubes above are $1110, 1111$, and so $111{*} \in F$. □

2.8.5 Codimension 4

Theorem 2.39 shows that an irreducible $(6,4)$-homogeneous subcube partition exists. Using techniques similar to the preceding section, in this section we show that an irreducible $(n,4)$-homogeneous subcube partition exists only for $n = 6$.

Theorem 2.42. *If there exists an irreducible* $(n,4)$*-homogeneous subcube partition then* $n = 6$*. Moreover, the irreducible* $(6,4)$*-homogeneous subcube partition is unique up to permutation and flipping of coordinates.*

Since the proof is a bit long, we break it into three parts, starting with the following lemma.

Lemma 2.43. *If* F *is an irreducible* $(n,4)$*-homogeneous subcube partition and* $s,t \in F$ *are two different subcubes, then* s,t *have at least two non-star coordinates in common.*

Proof. Since s,t conflict, they must have at least one non-star coordinate in common. We assume that they have exactly one non-star coordinate in common, and reach a contradiction.

Without loss of generality, s starts $0000{*}{*}{*}$ and t starts ${*}{*}{*}1000$. The subcube $u^{(1)}$ of F covering 10^{n-1} must start with 1 (to conflict with s) and its fourth symbol must be 0 (to conflict with t), and so it starts 1??0, where the question marks stand for 0 or $*$. Similarly, there are subcubes $u^{(2)}, u^{(3)}$ starting ?1?0 and ??10, respectively.

The subcubes $u^{(1)}, u^{(2)}$ must conflict, and so either $u^{(1)}$ starts $10{*}0$ or $u^{(2)}$ starts $01{*}0$ (the other question mark must be a star since otherwise $u^{(1)} \cup s$ or $u^{(2)} \cup s$ would be a subcube, contradicting irreducibility). Assume without loss of generality that $u^{(1)}$ starts $10{*}0$. Since $u^{(1)}, u^{(3)}$ must conflict, $u^{(3)}$ must start $0{*}10$. Since $u^{(2)}, u^{(3)}$ must conflict, $u^{(2)}$ must start ${*}100$.

The subcube $v \in F$ covering 1110^{n-3} must conflict with the subcube t which starts ${*}{*}{*}1000$, and so $v_4 = 0$. It must also conflict with $u^{(1)}$, and so $v_2 = 1$. Similarly $v_1 = v_3 = 1$, and so $v = 1110{*}^{n-4}$.

Repeating the same argument with subcubes covering $0001100\,0^{n-7}, 0001010\,0^{n-7}, 0001001\,0^{n-7}$, we obtain the following situation (without loss of generality), where we write only the first seven coordinates:

0000***	***1000
10*0???	???110*
*100???	???1*10
0*10???	???10*1
1110***	***1111

Every subcube $w \in F$ with $w_4 = 0$ other than $1110{*}^{n-4}$ must conflict with $1110{*}^{n-4}$, and so one of w_1, w_2, w_3 must be 0. Consequently, all such subcubes have weight at most 2.

Every subcube $w \in F$ with $w_4 = *$ must conflict with both $0000{*}^{n-4}, 1111{*}^{n-4}$, and so w_1, w_2, w_3 must contain both a 0 and a 1. Similarly, w_5, w_6, w_7 must contain both a 0 and a 1. Consequently, all such subcubes have weight exactly 2.

We conclude that all subcubes $w \in F$ of weight 3 other than $1110{*}^{n-4}$ must have $w_4 = 1$. According to Lemma 2.37, there are 3 of these (excluding $1110{*}^{n-4}$). The subcubes on the right side of the table above have weights $1, 2, 2, 2, 4$ (recall that question marks stand for 0 or $*$), and so there are at least 8 subcubes in F contained inside ${*}{*}{*}1{*}^{n-4}$. Considering sizes, there must be exactly 8 such subcubes, which together cover all of ${*}{*}{*}1{*}^{n-4}$, contradicting irreducibility. □

We use this observation together with the following lemma to bound n.

Lemma 2.44. *Let* F *be an irreducible* (n,k)*-homogeneous subcube partition, where* $n \geq 2$*. Each coordinate is mentioned in at least six subcubes of* F*.*

Proof. We will show that the first coordinate is mentioned at least six times. For $\sigma \in \{0,1,*\}$, let $F_\sigma = \{x : \sigma x \in F\}$. Since $F_0 \cup F_*$ and $F_1 \cup F_*$ are both subcube partitions, $\bigcup F_0 = \bigcup F_1$, and so $|F_0| = |F_1|$.

Lemma 2.11 shows that $|F_0| \geq 2$. Suppose that $|F_0| = 2$. The two subcubes in F_0 must conflict, and so without loss of generality, assume that $F_0 = \{0x, 1y\}$.

Recall that $|F_1| = 2$ and $\bigcup F_0 = \bigcup F_1$. If none of the subcubes in F_1 starts with a star then $F_1 = \{0x', 1y'\}$ (since $0x' \cup 0x''$ doesn't cover $1y$), and so $x' = x$ and $y' = y$. However, that implies that $00x, 10x \in F$, contradicting irreducibility.

Therefore at least one of the subcubes in F_1 starts with a star. If both start with a star then $\bigcup F_1$ starts with a star, implying that $x = y$ and so $00x, 01x \in F$, contradicting irreducibility. Hence without loss of generality, $F_1 = \{0x', *z'\}$.

Considering the points starting with 1, we see that $z' = y'$. However, this means that $01y', 1{*}y' \in F$, contradicting homogeneity. We conclude that $|F_0| = |F_1| \geq 3$. $\square$

We can now prove the theorem.

Proof of Theorem 2.42. Let F be an irreducible $(n, 4)$-homogeneous subcube partition. Suppose without loss of generality that $0000{*}^{n-4} \in F$. According to Lemma 2.43, every other subcube in F mentions at most two coordinates beyond the first four, and so at most $2 \cdot 15/6 = 5$ of these are mentioned at least six times. Lemma 2.44 implies that $n \leq 4 + 5 = 9$.

We can slightly improve on this, as follows. Let $u^{(1)}, u^{(2)}, u^{(3)}, u^{(4)}$ be the subcubes containing the points $1000\,0^{n-4}, 0100\,0^{n-4}, 0010\,0^{n-4}, 0001\,0^{n-4}$, respectively. Each of these subcubes must be different. Indeed, if for example $u^{(1)} = u^{(2)}$ then $u^{(1)} \supseteq 1000\,0^{n-4} \vee 0100\,0^{n-4} = {*}{*}00\,0^{n-4}$, which intersects with $0000{*}^{n-4}$.

Any two of $u^{(1)}, u^{(2)}, u^{(3)}, u^{(4)}$ must conflict, and so for distinct $i, j \in \{1, \ldots, 4\}$, either $u^{(i)}_j = 0$ or $u^{(j)}_i = 0$. This means that together, $u^{(1)}, u^{(2)}, u^{(3)}, u^{(4)}$ contain at least $\binom{4}{2} = 6$ zeroes among the first four coordinates. Therefore one of $u^{(1)}, u^{(2)}, u^{(3)}, u^{(4)}$ must contain at least $\lceil 6/4 \rceil = 2$ zeroes among the first two coordinates, and so mentions at most one coordinate beyond the first four.

This means that strictly fewer than $2 \cdot 15/6 = 5$ coordinates are mentioned at least six times, and so $n \leq 4 + 4 = 8$.

Recalling that $n \geq 5$ due to Lemma 2.34, we complete the proof of the theorem by checking with a computer that no irreducible $(n, 4)$-homogeneous subcube partitions exist for $n = 5, 7, 8$, and that the irreducible $(6, 4)$-homogeneous subcube partition is unique up to permutation and flipping of coordinates. $\square$

3 Nonbinary subcube partitions

So far we have considered subcube partitions of the hypercube $\{0, 1\}^n$. In this section, we study subcube partitions of $\{0, \ldots, q-1\}^n$ for arbitrary $q \geq 2$.

Definition 3.1 (Subcube partition)**.** A *subcube partition* of $\{0, \ldots, q-1\}^n$ (or: a subcube partition *over* $\{0, \ldots, q-1\}$ of length n) is a partition of $\{0, \ldots, q-1\}^n$ into *subcubes*, which are sets of the form

$$\{x \in \{0, \ldots, q-1\}^n : x_{i_1} = b_1, \ldots, x_{i_d} = b_d\}.$$

We identify subcubes with words over $\{0, \ldots, q-1, *\}$. The definitions of the following concepts are identical to the binary case: dimension and codimension of a subcube, point, edge, size (Definition 2.1); reducible subcube partition (Definition 2.2); tight subcube partition (Definition 2.3); conflicting subcubes (Definition 2.4).

Given a collection F of subcubes of $\{0, \ldots, q-1\}^n$, we can determine whether they form a subcube partition using the criterion of Lemma 2.6, replacing 2 with q. Determining whether a subcube partition of $\{0, \ldots, q-1\}^n$ is tight is easy using the definition, and we can determine irreducibility using Algorithm 1.

We start our exploration of subcube partitions over $\{0, \ldots, q-1\}$ in Section 3.1, where we show how to convert an irreducible subcube partition of $\{0, 1\}^n$ into an irreducible subcube partition of $\{0, \ldots, q-1\}^n$.

We then study the minimal size of tight irreducible subcube partitions over $\{0, \ldots, q-1\}$ in Section 3.2.

3.1 Expansion

In this section we show how to convert a subcube partition of $\{0, 1\}^n$ into a subcube partition of $\{0, \ldots, q-1\}^n$ in a way which preserves tightness and irreducibility.

Lemma 3.2. *Let F be a subcube partition of $\{0,1\}^n$, let $q \geq 2$, and let $\phi_1, \dots, \phi_n \colon \{0, \dots, q-1\} \to \{0,1\}$ be surjective functions.*

Extend the definitions of $\phi_1, \dots, \phi_n$ to $\{0, \dots, q-1, *\}$ by defining $\phi_i(*) = *$. Define a function $\phi \colon \{0, \dots, q-1, *\}^n \to \{0,1,*\}^n$ as follows: $\phi(\sigma_1 \dots \sigma_n) = \phi(\sigma_1) \dots \phi(\sigma_n)$. Let

$$G = \{s \in \{0, \dots, q-1, *\}^n : \phi(s) \in F\}.$$

Then

(a) G is a subcube partition of $\{0, \dots, q-1\}^n$.

(b) If F is tight then so is G.

(c) If F is irreducible then so is G.

Proof. We start by showing that G is a subcube partition. Notice first that the subcubes in G are disjoint. Indeed, suppose that $s, s' \in G$ are distinct. If $\phi(s) = \phi(s')$ then s, s' must disagree on a non-star position, and so conflict. If $\phi(s) \neq \phi(s')$ then $\phi(s), \phi(s')$ conflict at some position i, and s, s' conflict at the same position.

In order to show that the subcubes in G cover all of $\{0, \dots, q-1\}^n$, let $x \in \{0, \dots, q-1\}^n$. Since F is a subcube partition, $\phi(x)$ is covered by some subcube $t \in F$. Define a subcube s as follows: if $t_i = *$ then $s_i = *$, and otherwise $s_i = x_i$. Then $\phi(s) = t$ and so $s \in G$, and s covers x by definition.

Now suppose that F is tight. Then for every $i \in [n]$ there is a subcube $t \in F$ mentioning i. Since ϕ is surjective, we can find a subcube s mentioning i such that $\phi(s) = t$. Hence $s \in G$, and so G also contains a subcube mentioning i. Hence G is tight.

Finally, suppose that F is irreducible. If G is reducible then there is a a subset $H \subset G$, with $1 < |H| < |G|$, whose union is a subcube r. We claim that the union of $\phi(H) = \{\phi(s) : s \in H\}$ is the subcube $\phi(r)$.

Indeed, on the one hand, any $s \in H$ satisfies $s \subseteq r$ and so $\phi(s) \subseteq \phi(r)$, hence $\bigcup \phi(H) \subseteq \phi(r)$. On the other hand, let $x \in \phi(r)$ be an arbitrary point. Define a point $y \in \{0, \dots, q-1\}^n$ as follows: if $r_i = *$ then y_i is an arbitrary element of $\phi_i^{-1}(x_i)$, and otherwise $y_i = r_i$; in the latter case, $\phi_i(y_i) = \phi_i(r_i) = x_i$. By construction, $y \in r$, and so y is covered by some $s \in H$. Since $\phi(y) = x$, it follows that $\phi(s)$ covers x.

Since F is irreducible, either $|\phi(H)| = 1$ or $|\phi(H)| = |F|$. In the latter case, $\phi(r) = *^n$ and so $r = *^n$, implying that $H = G$, contrary to assumption. In the former case, $\phi(s) = \phi(r)$ for all $s \in H$. Choose two distinct subcubes $s, s' \in H$. Let $i \in [n]$ be a coordinate at which s, s' conflict. Since $r \supseteq s \vee s'$ we have $r_i = *$, and so $\phi(r)_i = *$. On the other hand, $s_i, s'_i \neq *$, contradicting $\phi(s) = \phi(s') = \phi(r)$. □

3.2 Minimal size

Section 2.3 studies the minimal size of a tight irreducible subcube partition of $\{0,1\}^n$. In this section we extend this study to tight irreducible subcube partitions of $\{0, \dots, q-1\}^n$, asking: what is the minimal size of a tight irreducible subcube partition of $\{0, \dots, q-1\}^n$?

Applying Lemma 3.2 to the tight irreducible subcube partitions constructed in Theorem 2.23, we obtain a tight irreducible subcube partition of size $(n-1)q(q-1)+1$. We conjecture that this is optimal.

Conjecture 7. *If $n \geq 3$ then for all $q \geq 2$, the minimal size of a tight irreducible subcube partition of $\{0, \dots, q-1\}^n$ is $(n-1)q(q-1)+1$.*

We formally describe the matching construction in Section 3.2.1, where we also show that this is the minimal size that can be achieved by a direct application of Lemma 3.2, assuming Conjecture 1.

We prove Conjecture 7 for $n = 3$ in Section 3.2.2, where we also show that no tight irreducible subcube partition exists for $n = 2$. We have also verified the conjecture using a computer for $n = 4$ and $q \leq 6$, as well as for $n = 5$ and $q = 3$.

We close the section by proving a modest lower bound of $(q-1)n+1$ on the size of a tight subcube partition of $\{0, \dots, q-1\}^n$, using the technique of Tarsi [AL86]. The lower bound applies more generally to tight minimal subcube covers, where it is sharp.

3.2.1 Construction

In this section we show how to construct tight irreducible subcube partitions of $\{0,\dots,q-1\}^n$ of size $(n-1)q(q-1)+1$ using Lemma 3.2, and explain why this is the minimal possible size when using the lemma, assuming Conjecture 1. We start with the construction.

Theorem 3.3. *For each $n \ge 3$ and $q \ge 2$ there exists a tight irreducible subcube partition of $\{0,\dots,q-1\}^n$ of size $(n-1)q(q-1)+1$.*

Proof. Theorem 2.23 constructs a tight irreducible subcube partition of $\{0,\dots,q-1\}^n$ whose weight vector is $1, n-1, n-1, 0, \dots, 0$. Applying Lemma 3.2 with the mappings ϕ_i given by $\phi_i(0) = 0$ and $\phi_i(1) = \cdots = \phi_i(q-1) = 1$ for all $i \in [n]$, we obtain a tight irreducible subcube partition of size

$$1 \cdot (q-1)^0 + (n-1) \cdot (q-1)^1 + (n-1) \cdot (q-1)^2 = (n-1)q(q-1)+1. \qquad \square$$

We now show that this construction is the optimal way of applying Lemma 3.2, assuming Conjecture 1.

Theorem 3.4. *Assume that Conjecture 1 holds for some $n \ge 3$. Let F be a tight irreducible subcube partition of $\{0,1\}^n$. Let G be a subcube partition obtained by an application of Lemma 3.2 on F, for some $q \ge 2$. Then G has size at least $(n-1)q(q-1)+1$.*

Proof. Let $g(z_1,\dots,z_n)$ be the size of G when Lemma 3.2 is applied with functions $\phi_1,\dots,\phi_n \colon \{0,\dots,q-1\} \to \{0,1\}$ such that $|\phi_i^{-1}(0)| = z_i$ for all $i \in [n]$. The function g is multilinear, and so its minimal value over $\{1,\dots,q-1\}^n$ is attained at some $z \in \{1,q-1\}^n$. Define a subcube partition F' by flipping all coordinates i such that $z_i = q-1$. Then

$$|G| \ge g(z) = \sum_{s \in F'} (q-1)^{\#_1(s)}.$$

Since F' is tight and irreducible, a combination of Theorem 2.22 and Lemma 2.20 shows that

$$|G| \ge \min\bigl(1 + (n-1)(q-1) + (n-1)(q-1)^2, 1 + n(q-1) + (n-3)(q-1)^2 + (q-1)^3.\bigr).$$

If we subtract the first sum from the second then we obtain

$$(q-1)^3 - 2(q-1)^2 + (q-1) = (q-2)^2(q-1) \ge 0,$$

and so the minimum equals the first sum. $\square$

3.2.2 Short length

In this section we characterize all tight irreducible subcube partitions of $\{0,\dots,q-1\}^n$ for $q \ge 2$ and $n \le 3$.

It is easy to see that the unique tight irreducible subcube partition of $\{0,\dots,q-1\}^1$ is $\{0,\dots,q-1\}$. In contrast, there is no tight irreducible subcube partition of $\{0,\dots,q-1\}^2$.

Lemma 3.5. *There are no tight irreducible subcube partitions of $\{0,\dots,q-1\}^2$ for any $q \ge 2$.*

Proof. Let F be a tight subcube partition of $\{0,\dots,q-1\}^2$. If all subcubes in F are points then F is clearly reducible. Otherwise, without loss of generality $0* \in F$. For every $a \in \{1,\dots,q-1\}$, let $F_a \subset F$ consist of all subcubes of F starting with a. Since F is tight, $F_a \neq \{a*\}$ for some a. Since $\bigcup F_a = a*$, it follows that F is reducible. $\square$

There is a unique tight irreducible subcube partition of $\{0,1\}^3$, up to flipping coordinates. An analogous result holds for all $q \ge 2$.

Lemma 3.6. *Every tight irreducible subcube partition of $\{0,\dots,q-1\}^3$, for any $q \ge 2$, can be obtained from $S_3 = \{000, 01*, 1*0, *01, 111\}$ by Lemma 3.2.*

Proof. Let G be a tight irreducible subcube partition of $\{0,\ldots,q-1\}^3$. Since G is tight, $*** \notin G$. Furthermore, no subcube in G contains two stars. Indeed, suppose that $0** \in G$. Then all subcubes in G starting with 1 cover $1**$, and all subcubes in G starting with 2 cover $2**$. Since G is irreducible, we see that $G = \{0**, 1**, 2**\}$, contradicting tightness.

Let $A(\cdot?*)$ denote the projection of all subcubes of G of the form $??*$ to the first coordinate, and define other A-sets analogously.

If $A(\cdot?*) = \emptyset$ then no subcube of G ends with $*$. Therefore the subcubes ending with $b \in \{0,1,2\}$ cover all of $**b$. Since G is irreducible, $**b \in G$, which is impossible. Therefore $A(\cdot?*) \neq \emptyset$.

We claim that $A(\cdot?*)$ and $A(\cdot*?)$ are disjoint. Indeed, if $a \in A(\cdot?*) \cap A(\cdot*?)$, then $ab*, a*c \in G$ for some $b, c \in \{0,1,2\}$, which is impossible since these subcubes intersect.

We claim that if $a \in A(\cdot?*)$ and $b \in A(?\cdot*)$ then $ab* \in G$. Indeed, suppose that $ab'*, a'b* \in G$ but $ab* \notin G$. Consider a point $abc \in \{0,\ldots,q-1\}^3$. This point cannot be covered by $a*c$ since this subcube does not conflict with $ab'*$, and cannot be covered by $*bc$ since this subcube does not conflict with $a'b*$. Therefore $abc \in G$. Since this holds for all c and $\bigcup_c abc = ab*$, we get a contradiction with the irreducibility of G.

It follows that G is composed of points and edges, where the edges are

$$\{ab* : a \in A(\cdot?*), b \in A(?\cdot*)\} \cup \{a*c : a \in A(\cdot*?), c \in A(?*\cdot)\} \cup \{*bc : b \in A(*\cdot?), c \in A(*?\cdot)\}.$$

We claim that $A(\cdot?*) \cup A(\cdot*?) = \{0,\ldots,q-1\}$, and so these two sets partition $\{0,\ldots,q-1\}$. Indeed, suppose that a is contained in neither set. Let $b \in A(?\cdot*)$, so that $b \notin A(*\cdot?)$. By construction, points of the form abc are not covered by any of the edges of G, hence all of them belong to G. Since $\bigcup_c abc = ab*$, this contradicts the irreducibility of G.

It follows that G can be obtained by applying Lemma 3.2 to S_3 with the mappings

$$\phi_1(a) = 1 \leftrightarrow a \in A(\cdot*?), \qquad \phi_2(b) = 1 \leftrightarrow b \in A(?\cdot*), \qquad \phi_3(c) = 1 \leftrightarrow c \in A(*?\cdot).$$

Indeed, the edges of G are

$$\{ab* : \phi_1(a) = 0, \phi_2(b) = 1\} \cup \{a*c : \phi_1(a) = 1, \phi_3(c) = 0\} \cup \{*bc : \phi_2(b) = 0, \phi_3(c) = 1\},$$

and these cover all points $abc \in \{0,\ldots,q-1\}^3$ other than the ones satisfying $\phi_1(a) = \phi_2(b) = \phi_2(c)$. ☐

Corollary 3.7. *Conjecture 7 holds for $n = 3$ and all $q \geq 2$.*

Proof. Let G be a tight irreducible subcube partition of $\{0,\ldots,q-1\}^3$, where $q \geq 2$. According to the lemma, it can be obtained by applying Lemma 3.2. The result now follows from Theorem 3.4, since it is known that all tight irreducible subcube partitions of $\{0,1\}^3$ have size 5. ☐

When $n \geq 4$, not all tight irreducible subcube partitions are obtained via Lemma 3.2. Here is an example:

$$0000, 0002, 0020, 0022, 0101, 0102, 0111, 0122, 0200, 0201, 0211, 0220, 1010, 1011, 1020, 1021, 1102,$$
$$1110, 1120, 1122, 1201, 1202, 1211, 1212, 2011, 2012, 2021, 2022, 2110, 2111, 2200, 2212, 2220, 2222,$$
$$2100, 2101, 01{*}0, 02{*}2, 0{*}21, 10{*}2, 11{*}1, 1{*}00, 20{*}0, 22{*}1, 2{*}02, \underline{001*}, \underline{122*}, \underline{212*}, \underline{*001}, *\underline{112}, \underline{*210}.$$

This is a tight irreducible subcube partition of $\{0,1,2\}^4$. The underlined subcubes show that it cannot be obtained by applying Lemma 3.2, since $001, 122, 212$ and $001, 112, 210$ are not product sets.

3.2.3 Lower bound

Theorem 2.9 gives our best lower bound on the size of a tight irreducible subcube partition of $\{0,1\}^n$, slightly improving on the "trivial" lower bound of $n+1$ which follows from the well-known lemma of Tarsi [AL86] on minimally unsatisfiable CNFs.

Tarsi's lemma applies more generally to subcube covers.

Definition 3.8 (Subcube cover)**.** A *subcube cover* of $\{0,\dots,q-1\}^n$ is a collection of subcubes whose union is $\{0,\dots,q-1\}^n$.

A subcube cover is *minimal* if no proper subset of it is a subcube cover.

In this language, Tarsi's lemma states that a tight minimal subcube cover of $\{0,1\}^n$ has size at least $n+1$. This bound is achieved, for example, by the subcube partition

$$\{0^i 1 *^{n-i-1} : 0 \le i \le n-1\} \cup \{0^n\}.$$

The analogous subcube partition for arbitrary $q \ge 2$ is

$$\{0^i b *^{n-i-1} : 0 \le i \le n-1, 1 \le b \le q-1\} \cup \{0^n\},$$

which has size $(q-1)n+1$.

In this section, we generalize Tarsi's lemma to the setting of matroids. A special case of our generalization shows that every tight minimal subcube cover of $\{0,\dots,q-1\}^n$ (and so every tight subcube partition of $\{0,\dots,q-1\}^n$) has size at least $(q-1)n+1$, proving the optimality of the above construction.

There are several proofs of Tarsi's lemma [AL86, CS88, ML97, DDKB98, Kul00, BET01]. We generalize the well-known proof using Hall's theorem.

Definition 3.9 (Cover)**.** Let M be a matroid. A collection F of subsets of the ground set of M is an *M-cover* if no basis of M intersects all sets in F. An M-cover is *minimal* if no proper subset is an M-cover.

Theorem 3.10 (Generalized Tarsi's lemma)**.** *Let M be a matroid with rank function r. Every minimal M-cover F satisfies*

$$|F| > r\left(\bigcup F\right).$$

The statement might look opaque, so before proving the theorem, we first show how it can be used to derive the lower bound $(q-1)n+1$.

Theorem 3.11. *Every tight minimal subcube cover of $\{0,\dots,q-1\}^n$, where $n \ge 1$ and $q \ge 2$, has size at least $(q-1)n+1$.*

Proof. Let $H(n,q)$ be the matroid over the ground set $[n] \times \{0,\dots,q-1\}$ in which a set is independent if for every $i \in [n]$, it doesn't contain all elements of the form $(i,?)$. A basis of $H(n,q)$ is any set of the form $B(a_1,\dots,a_n) := \{(i,j) : i \in [n], j \in [q], j \neq a_i\}$, where $a_1,\dots,a_n \in \{0,\dots,q-1\}$.

Let F be a tight minimal subcube cover of $\{0,\dots,q-1\}^n$. We can represent every subcube in $s \in F$ as the following subset of the ground set of $H(n,q)$:

$$\phi(s) = \{(i,s_i) : i \in [n], s_i \neq *\}.$$

Let $\phi(F) = \{\phi(s) : s \in F\}$. We claim that $\phi(F)$ is an $H(n,q)$-cover. Indeed, let $B(a_1,\dots,a_n)$ be any basis of $H(n,q)$. Since F is a subcube cover, the point $a_1 \dots a_n$ is covered by some subcube s. If $(i,s_i) \in \phi(s)$ then $s_i = a_i$, and so $\phi(s)$ is disjoint from $B(a_1,\dots,a_n)$.

A similar argument shows that $\phi(F)$ is a minimal $H(n,q)$-cover. Indeed, any proper subset of $\phi(F)$ has the form $\phi(G)$ for some proper subset $G \subset F$. Since F is a minimal subcube cover, some point $a_1 \dots a_n$ is not covered by G. The corresponding basis $B(a_1,\dots,a_n)$ intersects all sets in $\phi(G)$. Indeed, if $\phi(s) \in \phi(G)$ then s doesn't cover a, and so $s_i \neq a_i, *$ for some $i \in [n]$. Consequently, $\phi(s)$ contains $(i,s_i) \in B(a_1,\dots,a_n)$.

Since $\phi(F)$ is a minimal $H(n,q)$-cover, Theorem 3.10 shows that $|F| = |\phi(F)|$ exceeds the rank of $\bigcup \phi(s)$. We will show that $\bigcup \phi(s) = [n] \times \{0,\dots,q-1\}$, a set whose rank is $(q-1)n$, completing the proof.

Let $i \in [n]$. Since F is tight, some subcube $s \in F$ mentions i. Since F is minimal, there exists a point $x \in \{0,\dots,q-1\}^n$ which is only covered by s. In particular, no subcube of F contains $x^{i \to *}$, the subcube obtained from x by changing the i'th coordinate to a star. This implies that for every $b \in \{0,\dots,q-1\}$, every subcube of F containing $x^{i \to b}$ must contain b in its i'th coordinate. Therefore $\bigcup \phi(F)$ contains all elements of the form (i,b), for any $b \in \{0,\dots,q-1\}$, as promised. □

The proof of Theorem 3.10 uses a generalization of Hall's theorem to matroids.

Proposition 3.12 (Hall–Rado [Rad67, Wel71])**.** *Let M be a matroid with rank function r, and let F be a collection (multiset) of subsets of the ground set of M.*

If each subset $G \subseteq F$ satisfies $|G| \le r(\bigcup G)$ then we can choose an element e_s from each set $s \in F$ such that the elements e_s are distinct, and $\{e_s : s \in F\}$ is an independent set of M.

We can now prove Theorem 3.10.

Proof of Theorem 3.10. Let F be a minimal M-cover, and suppose that $|F| \le r(\bigcup F)$. We will show that this assumption leads to a contradiction.

If every subset $G \subseteq F$ satisfies $|G| \le r(\bigcup G)$ then Proposition 3.12 shows that F intersects the independent set $\{e_s : s \in F\}$. Since every independent set can be completed to a basis, this contradicts the assumption that F is an M-cover.

We conclude that some subset $G \subset F$ satisfies $|G| > r(\bigcup G)$. Among all such subsets, choose one which is inclusion-maximal. By assumption, $G \neq F$, and so by minimality, G is not an M-cover, say the basis B intersects all sets in G.

Let $M' = M/\bigcup G$ be the contraction of M by $\bigcup G$. The ground set of M' is the ground set of M with $\bigcup G$ removed, and its rank function is $r'(S') = r(S' \cup \bigcup G) - r(\bigcup G)$.

Let $F' = \{S \backslash \bigcup G : S \in F \backslash G\}$. Suppose that H' is a non-empty subset of F', say $H' = \{S \backslash \bigcup G : S \in H\}$. Since G is inclusion-maximal,

$$r'(H') = r\left(\bigcup H' \cup \bigcup G\right) - r\left(\bigcup G\right) = r\left(\bigcup (H \cup G)\right) - r\left(\bigcup G\right) > |H \cup G| - |G| = |H| = |H'|.$$

Applying Proposition 3.12, we obtain a basis B' of F' which intersects all sets in F', and so all sets in $F \setminus G$.

The set $B \cap \bigcup G$ is an independent subset of $\bigcup G$. Complete it to a basis B_G of $M|\bigcup G$. Since B' is a basis of M', $B' \cup B_G$ is a basis of M. By construction, B intersects all subsets in F, contradicting the assumption that F is an M-cover. □

4 Affine vector space partitions

Section 2 considers partitions of $\{0,1\}^n$ into subcubes. In this section, we consider partitions of $\{0,1\}^n$ into affine subspaces. The companion work [BFIK22] considers the more general case of partitions of $\mathbb{F}_q^n$ into affine subspaces.

Definition 4.1 (Affine vector space partition)**.** An *affine vector space partition* of length n is a partition of $\{0,1\}^n$ into *affine subspaces*, that is, sets of the form $x + V$, where $x \in \{0,1\}^n$ and V is a subspace of $\{0,1\}^n$ (identified with $\mathbb{F}_2^n$). The *size* of an affine vector space partition is the number of affine subspaces.

The *linear part* of an affine subspace $U = x + V$ is the subspace V. The *dimension* of an affine subspace is the dimension of its linear part, and codimension is defined analogously.

The notion of reducibility is defined just as in Definition 2.2 and Definition 2.12, replacing *subcube* with *affine subspace*.

Definition 4.2 (Reducibility)**.** A collection F of disjoint affine subspaces of $\{0,1\}^n$ is *reducible* if there exists a subset $G \subseteq F$, with $|G| > 1$, whose union is an affine subspace of $\{0,1\}^n$ other than $\{0,1\}^n$. If no such G exists then F is *irreducible*.

The definition of tightness is perhaps less obvious. A subcube partition of length n is not tight if it arises from a subcube partition of length $n-1$ via an embedding of $\{0,1\}^{n-1}$ inside $\{0,1\}^n$. If this is the case, then there is a direction i which is "ignored" by all subcubes, in the sense that $s_i = *$. This definition generalizes to our setting, where an affine subspace $x + V$ "ignores" a direction $y \in \{0,1\}^n \setminus \{0^n\}$ if $y \in V$.

Definition 4.3 (Tightness)**.** An affine vector space partition F of length n is *tight* if the intersection of the linear parts of all affine subspaces in F is $\{0^n\}$.

We can determine whether two affine subspaces intersect by solving linear equations. Using this, we can determine whether a collection of affine subspaces forms an affine vector space partition as in Lemma 2.6, by checking that

$$\sum_s 2^{-\operatorname{codim}(s)} = 1.$$

We can check tightness using the definition, and irreducibility using Algorithm 1, suitably generalized. For this we need to be able to compute the join of two affine subspaces, which is the minimal affine subspace containing their union.

Lemma 4.4. *The minimal affine subspace containing $a + V$ and $b + W$ is $a + \operatorname{span}(V, W, b - a)$.*

We leave the straightforward proof to the reader.

We commence the study of affine vector space partitions in Section 4.1, where we show how to convert an irreducible subcube partition to an irreducible affine vector space partition. We use this technique in Section 4.2 to construct tight irreducible affine vector space partitions of length n and size $\frac{3}{2}n - O(1)$. In the same section we also prove a lower bound of $n + 1$ on the size of a tight irreducible affine vector space partition of length n.

4.1 Compression

Every subcube partition of length n can be viewed as an affine vector space partition of length n. Furthermore, if the subcube partition is tight, then so is the affine vector space partition. However, irreducibility is not maintained in this conversion. For example,

$$000, 111, 01*, 1{*}0, {*}01$$

is irreducible as a subcube partition but reducible as an affine vector space partition, since $000 \cup 111$ is an affine subspace, which we can represent by aaa. If we merge these two points, we get the tight irreducible affine vector space partition

$$aaa, 01*, 1{*}0, {*}01.$$

In this section we generalize this process of merging for arbitrary irreducible subcube partitions.

Definition 4.5 (Star pattern)**.** The *star pattern* of a subcube $s \in \{0, 1, *\}^n$ is $P(s) := \{i \in [n] : s_i = *\}$.

Lemma 4.6. *Let F be an irreducible subcube partition. For $S \subseteq [n]$, let F_S consist of all subcubes in F whose star pattern is S. For each $S \subseteq [n]$, choose a partition of F_S in which the union of each part is an affine subspace, and let G_S be the corresponding collection of affine subspaces. (If $F_S = \emptyset$, take $G_S = \emptyset$.)*

If all G_S are irreducible then $G = \bigcup_S G_S$ is also irreducible.

Proof. If G is reducible then there exists a subset $G' \subset G$, with $|G'| > 1$, whose union is an affine subspace U other than $\{0, 1\}^n$. Each affine subspace in G' is a union of subcubes of F. Let F' be the collection of all such subcubes, so that $\bigcup F' = U$.

If F' contains a subcube s with $s_i = *$ then according to Lemma 4.4, the linear part of U contains $0^{i-1}10^{n-i}$. This motivates defining S as the set of coordinates $i \in [n]$ such that $s_i = *$ for some $s \in F'$. Note that $S \neq [n]$, since otherwise $U = \{0, 1\}^n$.

Let $U|_{\bar{S}}$ be the projection of U into the coordinates outside of S, so that

$$U = \{x \in \{0, 1\}^n : x|_{\bar{S}} \in U|_{\bar{S}}\}.$$

Let $y \in U|_{\bar{S}}$. Every $x \in \{0, 1\}^n$ such that $x|_{\bar{S}} = y$ is covered by some subcube $s \in F'$. The definition of S implies that $s_i = y_i$ for all $i \in \bar{S}$. Consequently the union of all subcubes $s \in F'$ such that $s|_{\bar{S}} = y$ is the subcube $s_y := \{x \in \{0, 1\}^n : x|_{\bar{S}} = y\}$. Since F is irreducible, $s_y \in F$ and so $s_y \in F'$. Since $s_y \in F_S$, it follows that $F' \subseteq F_S$, and so $G' \subseteq G_S$. This contradicts the irreducibility of G'. □

In general, F being tight doesn't guarantee that G is tight. For example, applying Lemma 4.6 to the tight subcube partition

$$*000, *111, 001*, 0*01, 01*0, 110*, 1*10, 10*1$$

results in the non-tight affine vector space partition

$$*aaa, aa\bar{a}*, a*a\bar{a}, a\bar{a}*a,$$

in which all linear parts contain the non-zero vector 1111.

The following lemma is a simplification of Lemma 4.6 which also includes a criterion for tightness.

Lemma 4.7. *Let F be an irreducible subcube partition. For $S \subseteq [n]$, let F_S consist of all subcubes in F whose star pattern is S.*

Suppose that whenever F_S is non-empty, the union of all subcubes in F_S is an affine subspace g_S (this is always the case when $|F_S| \leq 2$). Then $G = \{g_S : F_S \neq \emptyset\}$ is an irreducible affine vector space partition.

Furthermore, G is tight if

$$\bigcap_{S\colon F_S \neq \emptyset} P\left(\bigvee F_S\right) = \emptyset, \tag{1}$$

where the join is taken in the sense of subcubes.

Proof. The irreducibility of G follows directly from Lemma 4.6. Indeed, if for every non-empty F_S we take the partition consisting of a single part then the affine vector space partition G in this lemma coincides with that in Lemma 4.6. Moreover, if $F_S = \{a\}$ then a is itself an affine subspace, and if $F_S = \{a, b\}$ then $a \cup b$ is the affine subspace obtained from a by adding the following vector to the linear part: $v_i = 1$ if $a_i \neq b_i$ and $v_i = 0$ otherwise.

We proceed to show that if Equation (1) holds then G is tight. Let $S \subseteq [n]$ be such that F_S is non-empty. If $i \notin P(\bigvee F_S)$ then all $x \in g_S$ have the same value of x_i, and so $y_i = 0$ for all y in the linear part of g_S. Therefore if $y_i = 1$ for some y in the linear part of g_S then $i \in P(\bigvee F_S)$. Equation (1) thus guarantees that the only vector in the intersection of the linear parts of all g_S is the zero vector, and so G is tight. □

4.2 Minimal size

In Section 2.3 we conjectured that the minimal size of a tight irreducible subcube partition of length n is $2n - 1$. Using a computer, we have determined the minimal size of a tight irreducible affine vector space partition of length n for small n [BFIK22]:

n	3	4	5	6	7
scp	5	7	9	11	13
avsp	4	6	7	8	10

The first row is the minimal size of a tight irreducible subcube partition of length n, and the second row is the minimal size of a tight irreducible affine vector space partition of length n.

The constructions presented later in this section suggest the following conjecture.

Conjecture 8. *The minimal size of a tight irreducible affine vector space partition is $\frac{3}{2}n - o(n)$.*

We give a matching construction in Section 4.2.2. The best lower bound we are aware of is $n + 1$, which we prove in Section 4.2.1 using an argument similar to the proof of Theorem 3.10.

4.2.1 Lower bound

In this section, we adapt the proof of Theorem 3.10 to the setting of affine vector space partitions.

Theorem 4.8. *Every tight affine vector space partition of length $n \geq 1$ has size at least $n + 1$.*

As in Theorem 3.11, the lower bound holds more generally for every tight minimal affine vector space cover, a concept we do not define formally.

Proof. Let M be the matroid over $\{0,1\}^n$ in which a subset is independent if it is linearly independent, and let r be its rank function.

Suppose that F a tight irreducible affine vector space partition of length $n \geq 1$, and let $\mathbf{F} = \{V^\perp : x + V \in F\}$ (which we consider as a multiset).

We claim that $r(\mathbf{F}) = n$. Indeed, since F is tight,

$$\operatorname{span}(\{V^\perp : x + V \in F\}) = \left(\bigcap\{V : x + V \in F\}\right)^\perp = \{0^n\}^\perp = \{0,1\}^n,$$

and so $r(\mathbf{F}) = n$.

Suppose that every $\mathbf{G} \subseteq \mathbf{F}$ satisfies $|\mathbf{G}| \leq r(\bigcup \mathbf{G})$. According to Proposition 3.12, we can choose an element $y_{x+V} \in V^\perp$ for each $x + V \in F$ such that the elements y_{x+V} form an independent set. In particular, we can find an element $z \in \{0,1\}^n$ such that $\langle z, y_{x+V}\rangle \neq \langle x, y_{x+V}\rangle$ for all $x + V \in F$. By construction, $z \notin x + V$ for all $x + V \in F$, contradicting the assumption that F is an affine vector space partition.

It follows that there exists some subset $\mathbf{G} \subseteq \mathbf{F}$ satisfying $|\mathbf{G}| > r(\bigcup \mathbf{G})$. Among all such subsets, choose one which is inclusion-maximal. If $\mathbf{G} = \mathbf{F}$ then we are done, so suppose that $\mathbf{G} \neq \mathbf{F}$. Since $\mathbf{F}$ is an affine vector space partition, there is a point z_0 which is not covered by any subspace in $\mathbf{G}$.

Let $M' = M/\operatorname{span}(\mathbf{G})$, and let r' be its rank function. Let $\mathbf{F}' = \{V^\perp \setminus \operatorname{span}(\mathbf{G}) : V^\perp \in \mathbf{F} \setminus \mathbf{G}\}$. If $\mathbf{G}' \subseteq \mathbf{F}'$ then

$$r'(\mathbf{G}') = r(\mathbf{G}' \cup \operatorname{span}(\mathbf{G})) - r(\operatorname{span}(\mathbf{G})) = r(\mathbf{G}' \cup \mathbf{G}) - r(\mathbf{G}) > |\mathbf{G}' \cup \mathbf{G}| - |\mathbf{G}| = |\mathbf{G}'|,$$

using the inclusion-maximality of $\mathbf{G}'$. Hence Proposition 3.12 allows us to choose $y'_{x+V} \in V^\perp \setminus \operatorname{span}(\mathbf{G})$ for all $x + V \in F \setminus G$ such that these vectors are independent in M', which means that no linear combination of them lies in $\operatorname{span}(\mathbf{G})$ (and in particular, they are linearly independent).

Let z be a point such that $\langle z, y'_{x+V}\rangle \neq \langle x, y'_{x+V}\rangle$ for all $x + V \in F \setminus G$ and $\langle z, y\rangle = \langle z_0, y\rangle$ for all $y \in \operatorname{span}(\mathbf{G})$. By construction, z is not covered by any of the subspaces of $F \setminus G$. It is also not contained in any $x + V \in G$, since $x + V = \{w : \langle w, y\rangle = \langle x, y\rangle \text{ for all } y \in V^\perp\}$ and $\langle z, y\rangle = \langle z_0, y\rangle$ for all $y \in V^\perp$ (recalling that z_0 is not covered by G). This contradicts the assumption that F is an affine vector space partition. ☐

It is tempting to conjecture a common generalization of Theorem 3.11 and Theorem 4.8, namely that a tight affine vector space partition of $\mathbb{F}_q^n$ has size at least $(q-1)n+1$. Unlike Theorem 3.11 and Theorem 4.8, this cannot be true for tight minimal affine vector space covers for $q \geq 4$ (hence, the proof above cannot generalize): we can construct tight minimal affine vector space covers of size $(q-1)(n-3) + \frac{3}{2}(q+1) - 1$ for q odd, and we can construct tight minimal affine vector space covers of size $(q-1)(n-3) + q + q/p$ for $q = p^h$, where p is a prime. These constructions derive from the two examples of minimal blocking sets in a projective plane described in [BB86]. We leave the details for elsewhere.

4.2.2 Construction

In this section, we construct tight irreducible affine vector space partitions of length n and size $\frac{3}{2}n - O(1)$ for all $n \geq 3$, using Lemma 4.7. To construct the underlying subcube partitions, we use an inductive approach in the style of the constructions in Sections 2.3 and 2.4.

Lemma 4.9. *Let F, H be irreducible subcube partitions of length $n \geq 2$ satisfying the following conditions, where F_S consists of all subcubes in F whose star pattern in S:*

(i) $\{s \in F : s_1 = *\} = \{s \in H : s_1 = *\}$.

(ii) $|F_S|, |H_S| \leq 2$ *for all* S.

(iii) *Equation* (1) *holds for* H.

Let $m = m(H)$ *be the number of star patterns* S *such that* H_S *is non-empty, and let* $m' = m'(H)$ *be the number of those star patterns where* $1 \notin S$ *(that is, the first coordinate is not a star).*

For every $k \geq 0$ *there exists a tight irreducible affine vector space partition of length* $n + 2k$ *and size* $m + km'$.

Proof. For $N \geq 0$, let $F{*}^N = \{s{*}^N : s \in F\}$, and define ${*}^N F$ similarly.

Let $F^* = \{s \in \{0,1\}^{n-1} : {*}s \in F\}$. If F^* is empty then the union of the subcubes starting with $b \in \{0,1\}$ is $b{*}^{n-1}$. Since F is irreducible, $F = \{0{*}^{n-1}, 1{*}^{n-1}\}$. However, this contradicts Equation (1), using $n \geq 2$. Therefore F^* is non-empty.

We will construct an infinite sequence of subcube partitions $F^{(k)}$ such that the following hold:

(i) $F^{(k)}$ is an irreducible subcube partition of length $n + 2k$.

(ii) $F^*{*}^{2k+1} \subseteq F^{(k)}$.

(iii) Subcubes in $F^*{*}^{2k+1}$ have different star patterns from subcubes in $F^{(k)} \setminus F^*{*}^{2k+1}$.

(iv) $|F^{(k)}_S| \leq 2$ for all S.

(v) $m(F^{(k)}) = m + km'$.

(vi) Equation (1) holds for $F^{(k)}$.

The result then follows by applying Lemma 4.7 to $F^{(k)}$.

The starting point is $F^{(0)} = \{sa : as \in H, |a| = 1, |s| = n-1\}$. By assumption, $F^{(0)}$ is irreducible and $F^*{*} = H^*{*}$. By construction, $F^*{*} \subseteq F^{(0)}$, and all subcubes in $F^{(0)} \setminus F^*{*}$ end with a non-star. The remaining properties are by assumption.

Given $F^{(k)}$, we construct $F^{(k+1)}$ as follows. Apply Lemma 2.13 with $F_0 = {*}F^{(k)}$ and $F_1 = F{*}^{2k+1}$ to obtain a subcube partition $G^{(k)}$. We define $F^{(k+1)} = \{tab : abt \in G^{(k)}, |a| = |b| = 1, |t| = n + 2k\}$. Since $F^*{*}^{2k+1} \subseteq F^{(k)}$, we can explicitly write

$$F^{(k+1)} = \{t{*}{*} : t \in F^*{*}^{2k+1}\} \cup \{t0{*} : t \in F^{(k)} \setminus F^*{*}^{2k+1}\} \cup \{t1b : bt \in F{*}^{2k+1}, b \neq *\}.$$

We now verify the properties of $F^{(k+1)}$ one by one:

(i) By the induction hypothesis, $F^*{*}^{2k+1} \subseteq F^{(k)}$, and so ${*}F^*{*}^{2k+1} \subseteq {*}F^{(k)}$. Since F^* is non-empty, $F_0 \cap F_1 = {*}F^*{*}^{2k+1}$ is non-empty. Therefore Lemma 2.13 shows that $G^{(k)}$ is irreducible, and it follows that $F^{(k+1)}$ is irreducible.

(ii) The formula for $F^{(k+1)}$ immediately implies that $F^*{*}^{2k+3} \subseteq F^{(k+1)}$.

(iii) The formula for $F^{(k+1)}$ shows that all $s \in F^{(k+1)} \setminus F^*{*}^{2k+3}$ satisfy $s_{n+2k+1} \neq *$, and so have different star patterns from any subcube in $F^*{*}^{2k+3}$.

(iv) The subcubes in each of the three sets in the formula for $F^{(k+1)}$ have different star patterns. Since $|F^*_S| \leq 2$ for all S and $|F^{(k)}_S| \leq 2$ for all S, it follows that $|F^{(k+1)}_S| \leq 2$ for all S.

(v) Denote the three parts in the formula for $F^{(k+1)}$ by A, B, C. Clearly $m(A) = m(F^*)$. Since the star patterns of the subcubes in $F^*{*}^{2k+1}$ are different from the star patterns of the subcubes in $F^{(k)} \setminus F^*{*}^{2k+1}$, we have $m(B) = m(F^{(k)}) - m(F^*)$. Finally, $m(C) = m'$. We conclude that $m(F^{(k+1)}) = m(F^{(k)}) + m' = m + (k+1)m'$.

(vi) Since the star patterns of the subcubes in $F^{*}*^{2k+1}$ are different from the star patterns of the subcubes in $F^{(k)} \setminus F^{*}*^{2k+1}$, the induction hypothesis implies that the intersection of $P(\bigvee F_S)$ for all star patterns S appearing in $A \cup B$ is contained in $\{n+2k+1, n+2k+2\}$.

Equation (1) for F implies that some $s \in F$ satisfies $s_1 \neq *$, and so C is non-empty. All star patterns of subcubes in S do not contain $n+2k+1$ or $n+2k+2$, and so Equation (1) holds for $F^{(k+1)}$. ☐

Using this lemma, we construct tight irreducible affine vector space partitions of length n and size $\frac{3}{2}n - O(1)$ for all $n \geq 3$. Our construction matches the optimal values in the table appearing in the beginning of the section.

Theorem 4.10. *For all odd $n \geq 3$ there is a tight irreducible affine vector space partition of length n and size $\frac{3}{2}n - \frac{1}{2}$.*

There is a tight irreducible affine vector space partition of length 4 and size 6.

For all even $n \geq 6$ there is a tight irreducible affine vector space partition of length n and size $\frac{3}{2}n - 1$.

Proof. Consider the following subcube partitions:

$$S_3 = \{*01, 000, 111, 1*0, 01*\},$$
$$S_4 = \{*01*, 1000, 1111, 11*0, 1*01, 000*, 01**\},$$
$$T_6 = \{*0110*, *1101*, *001*1, *010*0, *00**0, *0*0*1, **111*, 011*0*, 110*1*, 010***, 11**0*\}.$$

Using Algorithm 1, one can check that they are irreducible. One checks directly that the prerequisites of Lemma 4.9 hold in all cases (with $H = F$).

Since $m(S_3) = 4$ and $m'(S_3) = 3$, Lemma 4.9 with $F = H = S_3$ constructs tight irreducible affine subspace partitions of length $3 + 2k$ and size $4 + 3k = \frac{3}{2}(3 + 2k) - \frac{1}{2}$.

Applying Lemma 4.7 directly to S_4, we obtain a tight irreducible affine subspace partition of length 4 and size 6.

Since $m(T_6) = 8$ and $m'(T_6) = 3$, Lemma 4.9 with $F = H = T_6$ constructs tight irreducible affine subspace partitions of length $6 + 2k$ and size $8 + 3k = \frac{3}{2}(6 + 2k) - 1$. ☐

Applying Lemma 4.9 with $F = S_3*$ and $H = S_4$ constructs tight irreducible affine subspace partitions of length $4 + 2k$ and size $6 + 3k = \frac{3}{2}(4 + 2k)$, which is slightly worse that what we get using $F = H = T_6$.

References

[AL86] Ron Aharoni and Nathan Linial. Minimal non-two-colorable hypergraphs and minimal unsatisfiable formulas. *J. Combin. Theory Ser. A*, 43(2):196–204, 1986.

[BB86] Aart Blokhuis and Andries Evert Brouwer. Blocking sets in desarguesian projective planes. *Bulletin of the London Mathematical Society*, 18(2):132–134, 1986.

[BET01] Sven Baumer, Juan Luis Esteban, and Jacobo Torán. Minimally unsatisfiable CNF formulas. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS*, 74:190–192, 2001.

[BFF90] Marc A. Berger, Alexander Felzenbaum, and Aviezri S. Fraenkel. Irreducible disjoint covering systems (with an application to Boolean algebra). *Discrete Appl. Math.*, 29(2-3):143–164, 1990. First International Colloquium on Pseudo-Boolean Optimization and Related Topics (Chexbres, 1987).

[BFIK22] John Bamberg, Yuval Filmus, Ferdinand Ihringer, and Sascha Kurz. Affine vector space partitions. Submitted, 2022.

[BOH90] Yigal Brandman, Alon Orlitsky, and John L. Hennessy. A spectral lower bound technique for the size of decision trees and two-level AND/OR circuits. *IEEE Transactions on Computers*, 39(2):282–287, feb 1990.

[CH11] Yves Crama and Peter L. Hammer. *Boolean functions*, volume 142 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2011. Theory, algorithms, and applications.

[CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *J. Assoc. Comput. Mach.*, 35(4):759–768, 1988.

[DD98] Gennady Davydov and Inna Davydova. Dividing formulas and polynomial classes for satisfiability. In *SAT'98, 2nd Workshop on the Satisfiability Problem*, pages 12–21, 1998.

[DDKB98] Gennady Davydov, Inna Davydova, and Hans Kleine Büning. An efficient algorithm for the minimal unsatisfiability problem for a subclass of CNF. *Ann. Math. Artificial Intelligence*, 23(3-4):229–245, 1998.

[FKW02] Ehud Friedgut, Jeff Kahn, and Avi Wigderson. Computing graph properties by randomized subcube partitions. In José D. P. Rolim and Salil Vadhan, editors, *Randomization and Approximation Techniques in Computer Science*, pages 105–113, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.

[For73] Rodney Forcade. Smallest maximal matchings in the graph of the d-dimensional cube. *J. Combinatorial Theory Ser. B*, 14:153–156, 1973.

[GK13] Matthew Gwynne and Oliver Kullmann. Towards a theory of good SAT representations. *CoRR*, abs/1302.4421, 2013.

[GPW18] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM J. Comput.*, 47(6):2435–2450, 2018.

[Iwa87] Kazuo Iwama. Complementary approaches to CNF Boolean equations. In *Discrete algorithms and complexity (Kyoto, 1986)*, volume 15 of *Perspect. Comput.*, pages 223–236. Academic Press, Boston, MA, 1987.

[Iwa89] Kazuo Iwama. CNF-satisfiability test by counting and polynomial average time. *SIAM J. Comput.*, 18(2):385–391, 1989.

[KB00] Hans Kleine Büning. On subclasses of minimal unsatisfiable formulas. *Discrete Appl. Math.*, 107(1-3):83–98, 2000. Boolean functions and related problems.

[Kis20] Andrzej P. Kisielewicz. On the structure of cube tiling codes. *European Journal of Combinatorics*, 89:103168, 2020.

[Kor84] Ivan Korec. Irreducible disjoint covering systems. *Acta Arith.*, 44(4):389–395, 1984.

[Kul00] Oliver Kullmann. An application of matroid theory to the SAT problem. In *15th Annual IEEE Conference on Computational Complexity (Florence, 2000)*, pages 116–124. IEEE Computer Soc., Los Alamitos, CA, 2000.

[Kul04] Oliver Kullmann. The combinatorics of conflicts between clauses. In Enrico Giunchiglia and Armando Tacchella, editors, *Theory and Applications of Satisfiability Testing*, pages 426–440, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[KZ13] Oliver Kullmann and Xishun Zhao. On Davis-Putnam reductions for minimally unsatisfiable clause-sets. *Theoret. Comput. Sci.*, 492:70–87, 2013.

[KZ16] Oliver Kullmann and Xishun Zhao. Unsatisfiable hitting clause-sets with three more clauses than variables, 2016.

[ML97] Shaohan Ma and Dongmin Liang. A polynomial-time algorithm for reducing the number of variables in MAX SAT problem. *Sci. China Ser. E*, 40(3):301–311, 1997.

[PS22] Tomáš Peitl and Stefan Szeider. Are hitting formulas hard for resolution?, 2022.

[Rad67] Richard Rado. Note on the transfinite case of Hall's theorem on representatives. *J. London Math. Soc.*, 42:321–324, 1967.

[Wel71] Dominic James Anthony Welsh. Generalized versions of Hall's theorem. *J. Combinatorial Theory Ser. B*, 10:95–101, 1971.