

SPANOIDS—AN ABSTRACTION OF SPANNING STRUCTURES, AND A BARRIER FOR LCCS*

ZEEV DVIR[†], SIVAKANTH GOPI[‡], YUZHOU GU[§], AND AVI WIGDERSON[¶]

Abstract. We introduce a simple logical inference structure we call a “spanoid” (generalizing the notion of a matroid), which captures well-studied problems in several areas. These include combinatorial geometry (point-line incidences), algebra (arrangements of hypersurfaces and ideals), statistical physics (bootstrap percolation), network theory (gossip/infection processes) and coding theory. We initiate a thorough investigation of spanoids, from computational and structural viewpoints, focusing on parameters relevant to the applications areas above and, in particular, to questions regarding locally correctable codes (LCCs). One central parameter we study is the “rank” of a spanoid, extending the rank of a matroid and related to the dimension of codes. This leads to one main application of our work, establishing the first known barrier to improving the nearly 20-year old bound of Katz–Trevisan (KT) on the dimension of LCCs. On the one hand, we prove that the KT bound (and its more recent refinements) holds for the much more general setting of spanoid rank. On the other hand we show that there exist (random) spanoids whose rank matches these bounds. Thus, to significantly improve the known bounds one must step out of the spanoid framework. Another parameter we explore is the “functional rank” of a spanoid, which captures the possibility of turning a given spanoid into an actual code. The question of the relationship between rank and functional rank is one of the main questions we raise as it may reveal new avenues for constructing new LCCs (perhaps even matching the KT bound). As a first step, we develop an entropy relaxation of functional rank to create a small constant gap and amplify it by tensoring to construct a spanoid whose functional rank is smaller than rank by a polynomial factor. This is evidence that the entropy method we develop can prove polynomially better bounds than KT-type methods on the dimension of LCCs. To facilitate the above results we also develop some basic structural results on spanoids including an equivalent formulation of spanoids as set systems and properties of spanoid products. We feel that given these initial findings and their motivations, the abstract study of spanoids merits further investigation. We leave plenty of concrete open problems and directions.

Key words. locally correctable codes, spanoids, entropy, percolation, matroid, union-closed family

AMS subject classifications. 68R01, 68Q01, 68P30, 94B99

DOI. 10.1137/19M124647X

1. Introduction. This (somewhat long) introduction will be organized as follows. We begin by discussing locally correctable codes (LCCs) and the main challenges they present as this was the primary motivation for this work. We proceed to define spanoids as an abstraction of LCCs, and state some results about their rank which hopefully illuminates the difficulties with LCCs in a new light. We continue by describing other natural settings in which the spanoid structure arises in the hope of motivating the questions raised in the context of LCCs and demonstrating their po-

*Received by the editors February 25, 2019; accepted for publication (in revised form) January 29, 2020; published electronically May 13, 2020. A preliminary version of our paper appeared in the conference *Proceedings of the 10th Innovations in Theoretical Computer Science (ITCS 2019)*.

<https://doi.org/10.1137/19M124647X>

Funding: The first and second authors were supported by NSF CAREER award DMS-1451191 and NSF grant CCF-1523816. The third author was supported by Jacobs Family Presidential Fellowship. The fourth author was partially supported by NSF grant CCF-1412958.

[†]Department of Computer Science and Department of Mathematics, Princeton University, Princeton, NJ 08540 (zeev.dvir@gmail.com).

[‡]Microsoft Research, Redmond, WA 98052 (sigopi@microsoft.com).

[§]Laboratory for Information and Decision Systems, MIT, Cambridge, MA 02139 (yuzhougu@mit.edu).

[¶]School of Mathematics, Institute for Advanced Study, Princeton, NJ 08540 (avi@math.ias.edu).

tential to contribute to research in other areas. We then turn to the investigation of functional rank of spanoids, which aims to convert them to actual LCCs. We conclude with describing some of the structural results about spanoids obtained here.

1.1. Locally correctable codes. The introduction of *locality* to coding theory has created a large body of research with wide-ranging applications and connections, from probabilistically checkable proofs, private information retrieval, program testing, fault-tolerant storage systems, and many others in computer science and mathematics. We will not survey these, and the reader may consult the surveys [Yek12, Dvi12]. Despite much progress, many basic questions regarding local testing, decoding and correcting of codes remain open. Here we focus on the efficiency of locally *correctable* codes, that we now define. Note that the related, locally *decodable* codes (LDCs), will not be discussed in this paper, as our framework is not relevant to them (LCCs can be converted to LDCs with a small loss in parameters).

DEFINITION 1.1 (*q-LCCs*). *A code $C \subseteq \Sigma^n$ is called a q -query LCC with error-tolerance $\delta > 0$ if for every $i \in [n]$ there is a family (called a q -matching) M_i , of at least δn disjoint q -subsets of $[n]$, with the following decodability property. For every codeword $c \in C$, and for every $i \in [n]$, the value of c_i is determined¹ by the values of c in coordinates S , for every q -subset S in M_i .²*

Intuitively, given a vector $c' \in \Sigma^n$ which results from corrupting less than (say) $\epsilon \delta n$ coordinates of a codeword $c \in C$, recovering c_i for any given $i \in [n]$ is simple. Picking a random q -subset from M_i and decoding c_i according to it will succeed with probability at least $1 - \epsilon$, as only an ϵ -fraction of these q -subsets can be corrupted.

In this paper we focus on the most well-studied and well-motivated regime where both the “query-complexity” q and the error-tolerance δ are constants. It is not hard to see that there are no LCCs with $q = 1$ (unless the dimension is constant) and so we will start with the first interesting case of $q = 2$. A canonical example of a 2-query LCC, which will serve us several times below, is the Hadamard code. Here $\Sigma = \mathbb{F}_2$. Let k be any integer and set $n = 2^k - 1$. Let A be the $k \times n$ matrix whose columns are all nonzero k -bit vectors. The Hadamard code $C_H \in \mathbb{F}_2^n$ is *generated* by A , namely H consists of all linear combinations of rows of A . Since every column of A can be written as a sum of (namely, *spanned* by) pairs of other columns in $(n-1)/2$ different ways, the matchings M_i suggest themselves, and so does the *linear* correcting procedure: Add the values in coordinates of the random pair S from M_i to determine the i th coordinate.

A central parameter of codes is their *rate*, capturing the redundancy between the *dimension*, namely the number of information bits encoded (here k), and the length of the codeword (here n). As in this paper this k will be a tiny function of n , we will focus on the *dimension* itself. Note that in the example above, as in every *linear* code, this dimension is also the *rank* of the generating matrix. In general codes, dimension may be fractional, and is defined as follows. All logarithms are in base 2 unless otherwise noted.

¹Through some function that does not depend on the codeword c .

²Our definition is a “zero-error” version of the standard definition. By “zero-error” we mean that for any codeword c , the value of c_i can be determined correctly (without error) from the coordinates of c at any q -subset S in the matching M_i . A more general definition would say that c_i can be computed from $c|_S$ with high probability, or even just slightly better than a random guess. Our definition is equivalent to the more general definition for linear codes, which comprise all of the interesting examples. We still allow “global” error in the sense that a large (constant) fraction of the coordinates can be corrupted (this global error tolerance is captured by the parameter δ).

DEFINITION 1.2 (dimension and rate of a code). *For a general, possibly non-linear code $C \subseteq \Sigma^n$, we define the dimension of C to be $\dim(C) = \log |C| / \log |\Sigma|$. Note that this coincides with the linear algebraic definition of dimension when C is a subspace. We refer to the ratio $\dim(C)/n$ as the “rate” of the code.*

Note that while the Hadamard code (C_H) has fantastic local correction (only two queries), its dimension is only $k \sim \log n$, which is quite bad from a coding theory perspective. However, no better 2-query LCC can exist, regardless of the alphabet.

THEOREM 1.3 (2-LCCs). *For all large enough n and over any alphabet,*

- *There exists a 2-query LCC of length n and dimension $\Omega(\log n)$ with some constant error tolerance δ (Folklore: Hadamard code).*
- *Every 2-query LCC of length n and error tolerance δ must have dimension at most $O_\delta(\log n)$ (for any constant δ) [BGT17].*

While we know precisely the optimal dimension for two queries, for $q \geq 3$ the gap between known upper and lower bounds is huge. The best lower bounds (constructions) are polylogarithmic: they come from Reed–Muller codes (using polynomials over finite fields), and yield dimension $\Omega((\log n)^{q-1})$.

The best LCC upper bounds are only slightly sublinear, giving $\dim(C) \leq \tilde{O}(n^{1-\frac{1}{q-1}})$ (up to logarithmic factors). This bound, which we will refer to as the Katz–Trevisan (KT) bound, is actually a slight refinement/improvement over the bound originally appearing in [KT00] (which gave $n^{1-1/q}$). This improvement was implicit in several works (e.g., [DK11, Woo07]) and is explicitly stated in [IS18]. We should also note that, over constant-size alphabets, Kerenidis and De Wolf proved an even stronger bound using quantum information theory [KW04].³ This exponential gap between the upper and lower bounds, which we formally state below, has not been narrowed in over two decades.⁴ Explaining this gap (in the hope of finding ways to close it) is one major motivation of this work.

THEOREM 1.4 (q -LCCs, $q \geq 3$). *For every fixed $q \geq 3$ and all large enough n ,*

- *There exists a q -query LCC of length n and dimension $\Omega((\log n)^{q-1})$ (with some constant tolerance δ and alphabet of size $q+1$) (Reed–Muller codes; see, e.g., the survey [Yek12]).*
- *Every q -query LCC of length n and error tolerance δ must have dimension at most $\tilde{O}_\delta(n^{1-\frac{1}{q-1}})$ (for any constant δ and any alphabet) [IS18].*

1.2. Spanoids. We shall now abstract the notion of inference used in LCCs. There, for a collection of pairs (T, i) with $T \subseteq [n]$ and $i \in [n]$, the values of codewords in coordinate positions T , determine the value of some other coordinate i . We shall forget (for now) the underlying code altogether, and abstract this relation by the formal “inference” symbol $T \rightarrow i$, to be read “ T spans i .”

DEFINITION 1.5 (spanoid). *A spanoid \mathcal{S} over $[n]$ is a family of pairs (T, i) with $T \subseteq [n]$ and $i \in [n]$. The pair (T, i) will sometimes be written as $T \rightarrow i$ and read as*

³Over constant size alphabet, the best known upper bound on the dimension of a q -query LCC (also LDC) is $\tilde{O}_\delta(n^{1-\frac{1}{\lceil q/2 \rceil}})$ [KW04]. This is better than the upper bound in Theorem 1.4 for $q \geq 4$. But for growing alphabet, Theorem 1.4 gives the best known upper bound.

⁴For LDCs, better constructions than Reed–Muller codes are known, through the seminal works of [Yek08, Efr09]. They construct 3-query LDCs of dimension $(\log n)^{\tilde{O}(\log \log n)}$. Still, the upper bounds for LDCs are the same as for LCCs, and obtained by the same KT-type argument, so the results in this paper may serve to better understand the (smaller, but still quite large) gap between upper and lower bounds in LDCs as well.

T spans i in the spanoid \mathcal{S} .

One natural way to view a spanoid is as a logical inference system, with the pairs indicating all inference rules. The elements of $[n]$ indicate some n formal statements, and an inference $T \rightarrow i$ of the spanoid means that if we know the truth of the statements in T , we can infer the truth of the i th statement. With this intuition, we shall adopt the convention that the inferences $i \rightarrow i$ are implicit in any spanoid, and that *monotonicity* holds: If $T \rightarrow i$, then also $T' \rightarrow i$ for every $T' \supseteq T$. These conventions will be formally stated below when we define general derivations, which sequentially combine these implicit rules and the stated rules (pairs) of the spanoid.

A key concept of spanoids is, naturally, the *span*. Given a subset $T \subseteq [n]$ (which we can think of as “axioms”), we can explore everything they can span by a sequence of applications of the inference rules of the spanoid \mathcal{S} .

DEFINITION 1.6 (derivation, span). *A derivation in \mathcal{S} of $i \in [n]$ from $T \subseteq [n]$, written $T \models_{\mathcal{S}} i$, is a sequence of sets $T = T_0, T_1, \dots, T_r$ with $i \in T_r$ such that for each $j \in [r]$, $T_j = T_{j-1} \cup i_j$ for some $i_j \in [n]$ and there exists $A \subset T_{j-1}$ such that $(A, i_j) \in \mathcal{S}$ is one of the spanoid rules.*

The span (or closure) of T , denoted $\text{span}_{\mathcal{S}}(T)$, is the set of all i for which $T \models_{\mathcal{S}} i$. We shall remove the subscript \mathcal{S} from these notations when no confusion about the underlying spanoid can arise, and write $T \models i$ and $\text{span}(T)$ for short.

Despite being highly abstract, we will see that spanoids can lead to a rich family of questions and definitions. The first, and perhaps one of the most central definitions, is that of the *rank* of a spanoid. We shall see other notions of spanoid rank later on (and will discuss the relation between them).

DEFINITION 1.7 (rank). *The rank of a spanoid \mathcal{S} , denoted $\text{rank}(\mathcal{S})$, is the size of the smallest subset $T \subseteq [n]$ such that $\text{span}(T) = [n]$. Note that by the definition of span we always have $\text{rank}(\mathcal{S}) \leq n$.*

We note that the “rank” of a logical inference system does appear (under different names) in *proof complexity*. It is the starting point for *expansion-based* lower bounds on a variety of proof systems, as introduced for resolution proofs in [BW01], and used for many others, e.g., in [ABRW04] and [AR01]). We shall return to this connection presently.

We can now define the spanoid analog of q -LCCs as spanoids which only specify the correction structure (the matchings M_i) without requiring any codewords or alphabet.

DEFINITION 1.8 (q -LCS, locally correctable spanoid). *A spanoid \mathcal{S} over $[n]$ is a q -LCS with error-tolerance δ if for every $i \in [n]$ there exists a family M_i of at least δn disjoint q -subsets of $[n]$ such that for each $S \in M_i$ we have $(S, i) \in \mathcal{S}$. Namely, each $i \in [n]$ is spanned (in \mathcal{S}) by at least δn disjoint subsets of q -elements.*

One can now ask about the highest possible rank of a q -LCS. It is not hard to see that the existence of a q -LCC (over any alphabet) $C \subset \Sigma^n$ with dimension $\dim(C) = d$ automatically implies that there exists a q -LCS (namely, the one given by the same matchings used in C) with rank at least $\lceil d \rceil$. Indeed, otherwise there would be $r < d$ coordinates in $[n]$ that determine any codeword $c \in C$ and this would limit the number of codewords to Σ^r .

One of our main observations is that, remarkably, in locally correctable spanoids there is *no* gap between the upper and lower bounds: We know the precise answer up to logarithmic factors, and it matches the *upper bounds* for LCCs! Observe the

analogies to the theorems in the previous subsection, for $q = 2$ and $q \geq 3$.

THEOREM 1.9 (2-LCSs). *For all large enough n ,*

- *There exists a 2-LCS over $[n]$ with error-tolerance δ of rank $\Omega(\frac{1}{\delta} \log(\delta n))$.*
- *Every 2-LCS over $[n]$ with error-tolerance δ must have rank at most $O(\frac{1}{\delta} \log(n))$.*⁵

Here, of course, the inference structure of the Hadamard code proves the first item. To get the required dependence on δ , one can take $\frac{1}{\delta}$ disjoint copies of such spanoids. The second item requires a new proof we discuss below, which generalizes (and implies) the one in Theorem 1.3. It is quite surprising that, even in this abstract setting, with no need for codewords or alphabet, one cannot do better than the Hadamard code!

We now state our results for $q \geq 3$.

THEOREM 1.10 (q -LCSs with $q \geq 3$). *For every fixed $q \geq 3$ and all large enough n ,*

- *There exist a q -LCS of rank $\tilde{\Omega}(n^{1-\frac{1}{q-1}})$ (with constant δ).*
- *Every q -LCS over $[n]$ has rank at most $\tilde{O}_\delta(n^{1-\frac{1}{q-1}})$ (for any constant δ).*

Both parts of this theorem demand discussion. The possibly surprising (and tight) lower bound follows from a simple probabilistic argument (indeed, one which is repeatedly used to prove expansion in the proof complexity references cited above), where the matchings M_i are simply chosen uniformly at random. It seems to reveal how significant a relaxation spanoids are of LCCs (where probabilistic arguments fail completely). However, the best known LCC upper bound (Theorem 1.4) does not rule out the possibility that, at least for large alphabets, the two (LCC's dimension and LCS's rank) have the same behavior! From a more pessimistic (and perhaps more realistic) perspective, our lower bound shows the limitations of any (upper bound) proof technique which, in effect, applies also for spanoids. These are proofs in which the LCC structure is used to show that a small subset spans all the others. We note that there are several LCC upper bounds which “beat” the $n^{1-\frac{1}{q-1}}$ bound for certain very special cases by using additional structure not present in the corresponding abstract spanoid. One example is the bound of [KW04], which uses arguments from quantum information theory to roughly *halve* the number of queries, over binary (or small) alphabets. Another example is the paper [DSW14], which gives an improved upper bound on the dimension for linear 3-LCCs defined over the real numbers, using specific properties of the Reals such as distance and volume arguments.

Our proof of the upper bound, is again more general than for LCCs, and interesting in its own right. We use a simple technique which performs random restrictions and contractions of graphs and hypergraphs (and originates in [DSW14]). It will be described in section 3, after we have formulated an equivalent, set-theoretic formulation of spanoids in section 2.2. Though the underlying idea of using random restrictions to prove our upper bounds for q -LCSs is similar to prior work on upper bounds for q -LCCs, our proofs are simpler and give better dependence on δ . In other words, we are proving more general theorems with simpler proofs. For example, the $O_\delta(\log n)$ upper bound for 2-LCCs in [BGT17] is proved by showing that there is a subset of $O_\delta(\log n)$ coordinates from which can infer the entire codeword. Our proof of Theorem 1.9 also shows this, but in a much more direct and simpler way. The proof of $\tilde{O}_\delta(n^{1-\frac{1}{q-1}})$ upper bound for q -LCCs in [IS18] is proved by a reduction to

⁵[BGT17] prove an upper bound of $O(\text{poly}(1/\delta) \log(n))$ on the rank of a 2-LCC with error-tolerance δ . But their proof also works for 2-LCSs.

the 2-LCC upper bound (using random restrictions). The article [IS18] also proves the $O_\delta(\log n)$ upper bound for 2-LCCs in a completely different way, but their proof only applies to linear 2-LCCs.

1.2.1. Functional rank: Bridging the gap between LCCs and LCSs. We conclude this section of the introduction with an attempt to understand (and possibly bridge) the gap between LCCs and their spanoid abstraction. The idea is to start with an LCS of high rank (which we know is possible), and convert it to an LCC without losing too much in the parameters. More generally, for a given spanoid \mathcal{S} , we would like to investigate the code C with largest dimension (over any alphabet Σ) which would be consistent with the inferences of \mathcal{S} . This is captured in the notion of *functional rank* which we now define.

DEFINITION 1.11 (functional rank). *Let \mathcal{S} be a spanoid over $[n]$. A code $C \subset \Sigma^n$ is consistent with \mathcal{S} if for every inference (S, i) in \mathcal{S} , and for every codeword $c \in C$, its values of coordinates S determine its value in coordinate i (by some fixed function, $f_{S,i}$ not depending on c).⁶*

Define the functional rank of \mathcal{S} , denoted $f\text{-rank}(\mathcal{S})$, to be equal to the supremum of the dimension $\dim(C)$, over all possible finite alphabets Σ and codes $C \subset \Sigma^n$ which are consistent with \mathcal{S} .

Of course, the strategy of constructing LCCs in two stages as above can only work if we can bound the gap between $\text{rank}(\mathcal{S})$ and $f\text{-rank}(\mathcal{S})$. This question, of bounding this gap or proving it can be large, is perhaps the most interesting one we raise (and leave mostly open for now). For now, we are able to show an example in which the two are different. The example providing a gap is depicted in Figure 1, arranging the coordinates as the vertices of a pentagon, the pair of vertices of each edge span the vertex opposite to it. That is, $\{x_1, x_2\} \rightarrow x_4, \{x_2, x_3\} \rightarrow x_5$, etc.




FIG. 1. The pentagon spanoid Π_5 where each coordinate is spanned by the coordinates of the opposite edge.

THEOREM 1.12 (constant gap between rank and functional rank). *The pentagon spanoid Π_5 depicted in Figure 1 has $\text{rank}(\Pi_5) = 3$ but $f\text{-rank}(\Pi_5) = 2.5$.*

Seeing that $\text{rank}(\Pi_5) = 3$ is easy by inspection. The lower bound of 2.5 on functional rank comes from a set-theoretic construction of consistent codes we will describe (in more generality) in section 5.2, where we develop a linear programming (LP) relaxation for $\text{rank}(\mathcal{S})$ called $\text{LP}^{\text{cover}}(\mathcal{S})$, but surprisingly this LP captures the

⁶One can think of a code consistent with \mathcal{S} also as a “representation” of \mathcal{S} in the spirit of matroid theory.

best set-theoretic construction of consistent codes. But even in this small example, the upper bound on functional rank is nontrivial to determine, as we allow all possible alphabets and consistent codes. Not surprisingly, Shannon entropy is the key to proving such a bound. In section 5.1, we develop an LP relaxation, based on entropy whose optimum $\text{LP}^{\text{entropy}}(\mathcal{S})$ upper bounds $\text{f-rank}(\mathcal{S})$. In this example, it proves 2.5 is the optimum.

One natural way of amplifying gaps as in the example above, which may also be useful in creating codes of high functional rank, is the idea of *tensoring*. In section 6, we develop different notions of tensoring *spanoids* inspired by tensoring of codes. In particular, we define a product of spanoids called the *semi-direct product* under which rank is multiplicative and f-rank is sub-multiplicative. By repeatedly applying this product to Π_5 , we get a spanoid with polynomial gap between f-rank and rank .

THEOREM 1.13 (polynomial gap between rank and function rank). *There exists a spanoid \mathcal{S} on n elements with $\text{rank}(\mathcal{S}) \geq n^c \text{f-rank}(\mathcal{S})$ where $c = \log_5 3 - \log_5 2.5 \geq 0.113$.*

Summarizing, we have the following obvious inequalities between the measures we described so far for every spanoid \mathcal{S} . We feel that understanding the exact relationships better is worthy of further study

$$(1.1) \quad \text{LP}^{\text{cover}}(\mathcal{S}) \leq \text{f-rank}(\mathcal{S}) \leq \text{LP}^{\text{entropy}}(\mathcal{S}) \leq \text{rank}(\mathcal{S}).$$

1.3. Other motivations and incarnations of spanoids. We return to discuss other structures: combinatorial, geometric, and algebraic, in which the same notions of span and inference naturally occur, leading to a set-theoretic one that elegantly captures spanoids precisely. These raise further issues, some of which we study in this paper and some are left for future work. These serve to illustrate the breadth of the spanoid framework.

1.3.1. Bootstrap percolation and gossip processes. The following general set-up occurs in statistical physics, network theory, and probability theory. Fix an undirected graph $G([n], E)$. In a gossip or infection process, or equivalently bootstrap percolation, we are given a set of “rules” specifying, for every vertex $v \in [n]$, a family of subsets of its neighbors. The intended meaning of such a rule is that if *every* member of one such subset is “infected” at a certain time step, then the vertex v becomes infected in the next time step. Given a set of initial infected vertices, this defines a process in which infection spreads, and eventually stabilizes. A well-studied special case is the (uniform) *r-bond percolation* [Bol68], where the family for each vertex is all r -subsets of its neighbors. Many variants exist, e.g., one can have a similar process on the edges, rather than vertices of the graph. An important parameter of such a process is the following: What is the size of the smallest set of vertices which, if infected, will eventually infect all other vertices.⁷

A moment’s thought will convince the reader that this structure is precisely a spanoid (where inferring sets are restricted by the graph structure). The infection process is precisely the inference process defining span in spanoids. Furthermore, the smallest size of an infecting set is *precisely* the rank of that spanoid! Much work has been invested to determine that rank even in very special cases, e.g., for the r -bond percolation above, in, e.g., Boolean hypercubes, where it is known precisely. Interestingly, the paper [HHQ17] uses the so-called “polynomial method” to reprove that

⁷This turns out to be crucial for understanding, at least for certain structured graphs like lattices studied by physicists, the threshold probability for percolation when initial infections are random.

bound, which fits even deeper with our framework. In our language, their method determines the *functional rank* of this spanoid, and one direction is through constructing an explicit code that is consistent with the spanoid! The reader is encouraged to work out the details.

1.3.2. Independence systems and matroids. An *independence system* over $[n]$ is a family \mathcal{F} of subsets of $[n]$ which is downwards-closed (if a set is in \mathcal{F} , so are all its subsets). The members of \mathcal{F} are called *independent*. While much of what we say below generalizes to all independence systems, we specify them for the important special systems called *matroids*.

A *matroid* is an independence system in which the independent sets satisfy the so-called “exchange axiom” (which we will not define here). Matroids abstract *linear* independence in subsets of a vector space over a field,⁸ and capture algorithmic problems in which optimization is possible through the greedy algorithm. Matroids thus come with natural notions of span and rank, extending the ones in the linear algebraic setting. The rank of a set is the size of the largest independent set it contains. The span of a set is the maximal superset of it of the same rank. A matroid can thus be naturally viewed as a spanoid, with the inference rules $F \rightarrow i$ for every independent $F \in \mathcal{F}$ and every i for which $F \cup \{i\}$ is *not* independent (such minimal dependent sets as $F \cup \{i\}$ are called *cycles*). It is easy to verify that the notions of span and rank of the matroid and the spanoid it defines coincide. This also raises the natural question of bounding the gap between *f-rank* and *rank* for the special case of spanoids arising from matroids.

Note that a spanoid resulting from a matroid this way is *symmetric*: By the exchange property of matroids, if $E \subset [n]$ is a cycle of \mathcal{F} , then for *every* $i \in E$ it contains the inference $E \setminus \{i\} \rightarrow i$. Symmetric spanoids are interesting, and we note that the pentagon example witnessing the gap between rank and functional rank is *not* symmetric, and we do not know such a gap for symmetric spanoids. We also don’t know if symmetric spanoids can achieve the lower bound in Theorem 1.10.

1.3.3. Point-line incidences. Sylvester–Gallai theorem is a celebrated result in combinatorial geometry conjectured by Sylvester and proved independently by Melchior and Gallai. It states that for any set of n points in Euclidean space \mathbb{R}^d , if the line through any two points passes through a third point, then they must all be collinear (namely, they span a 1-dimensional affine space). Over the complex numbers, one can prove a similar theorem but with the conclusion that the points span a 2-dimensional affine space (and there are in fact 2-dimensional examples known) [Kel86]. Over finite fields the conclusion is even weaker, saying that the span has dimension at most $O(\log(n))$ [BDSS11] and this is tight as the example of all points in \mathbb{F}_p^k with $n = p^k$ shows. It is not a coincidence that this example reminds one of the Hadamard code described before as an example of a 2-query LCC. It is in fact true that there is a tight connection between configurations of points with many collinear triples and linear 2-query LCCs. This was first noticed in [BDYW11, BDWY12] and was used to prove that 2-LCCs do not exist over the characteristic zero fields (for $q \geq 3$ these questions are wide open with even larger gaps than in the finite field case). LCCs with more than two queries naturally correspond to point configurations with many $(q - 1)$ -dimensional affine spaces containing at least $q + 1$ points.

Given the connections between Sylvester–Gallai type incidence structures and

⁸Matroids are in fact more general than linear independent sets of vectors over a field; for example, the Vámos matroid on eight elements is not representable over any field.

LCCs, and the insights offered by spanoids for studying LCCs, it is natural that the study of the spanoid structures can help us get new insights on incidence geometry problems. A *dual* way to view these incidences, which we shall presently generalize, is to consider each point $p_i : i \in [n]$ as representing a hyperplane F_i through the origin (in the appropriate vector space) vanishing on the linear function defined by p_i . A point p_i is spanned by a collection of points $\{p_j : j \in S\}$ iff $F_i \supset \bigcap_{j \in S} F_j$. Therefore, the spanning structure of the points p_1, p_2, \dots, p_n is captured by the spanoid where we would add the inference $S \rightarrow i$ iff F_i contains the common intersection of all $F_j : j \in S$.

1.3.4. Systems of polynomial equations. Given the above example, there is no reason to stop at the linear setting. Instead of lines we can consider n (multivariate) polynomials f_i over a field, and again consider F_i to be the zero set of f_i . The spanoid above, having an inference $S \rightarrow i$ whenever the set F_i contains the common intersection of all $F_j : j \in S$, is capturing another natural algebraic notion. Namely, it says that the polynomial f_i vanishes on all of the common roots of the polynomials $\{f_j : j \in S\}$. By the celebrated Hilbert’s Nullstellensatz theorem, over algebraically closed fields, this implies that f_i belongs to the radical of the ideal generated by the f_j ’s. Here the rank function is far from being that of a matroid; the complex spanoid which arises (and in general is far from understood) plays a role in arithmetic complexity (a beautiful example is the recent [Shp18] dealing with degree-2 polynomials).

1.3.5. Intersecting set systems. Let us remove all restrictions from the origin or nature of the n sets F_i discussed in the previous discussion. Assume we are given any such family \mathcal{F} of sets (from an arbitrary universe, say U). As above, a natural spanoid $\mathcal{S}_{\mathcal{F}}$ will have the inference $S \rightarrow i$ whenever the set F_i contains the common intersection of all $F_j : j \in S$. Such situations (and hence, spanoids) arise in many questions of extremal set theory; for example, the study of (weak) sunflowers, or families with certain forbidden intersection (or union) patterns, e.g., [FLS12, EFF85, F96].

What is interesting in this more general framework, where the initial family of sets \mathcal{F} is arbitrary, is that it becomes *equivalent* to spanoids! In other words, every spanoid \mathcal{S} arises as $\mathcal{S}_{\mathcal{F}}$ of *some* family of sets \mathcal{F} . This possibly surprising fact is not much more than an observation, but it turns out to be an extremely useful formulation for proving some of the results in this paper. Let us state it formally (it will be proved in section 2.2).

THEOREM 1.14 (spanoids and intersecting sets). *Let \mathcal{S} be any spanoid on $[n]$. Then, there exists a universe U and a family \mathcal{F} of n sets of U , $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$, such that $\mathcal{S} = \mathcal{S}_{\mathcal{F}}$.*

It is convenient to assume that the sets in \mathcal{F} have no element in common to all.⁹

The notions of rank and span are extremely simple in this set-theoretic setting, and do not require the sequential “derivation” and the implicit ordering which we require to define these in spanoids. For a family \mathcal{F} of n sets and a subset $S \subset [n]$, let us denote by $\bigcap S$ the subset of U which is the intersection of all $\{F_j : j \in S\}$. Then the rank of S is the size of the smallest subset $S' \subseteq S$ for which $\bigcap S' = \bigcap S$. Similarly, the span of S is the largest superset $S'' \supseteq S$ for which $\bigcap S'' = \bigcap S$.

These static definitions of rank and span make many things transparent. For example, the expected fact that testing if the rank of a spanoid (namely the rank of the

⁹Indeed, otherwise we can remove the common intersection (if nonempty) of *all* members of \mathcal{F} from each of them, as it does not change the underlying spanoid.

set $[n]$ is at most some given integer k is *NP-complete* (Claim 2.8). Complementing the sets F_i in \mathcal{F} , and replacing intersection with union, is precisely the set cover problem. This connection also underlies the cover-based linear program discussed earlier, as well as proofs of the main quantitative results, Theorems 1.9 and 1.10.

1.3.6. Union-closed families. Spanoids over $[n]$ are equivalent to union-closed families of subsets of $[n]$ i.e., a family of subsets of $[n]$ such that the union of any two members is again in the family. A closed set of a spanoid \mathcal{S} is a subset $A \subset [n]$ such that $\text{span}(A) = A$. The family of all closed sets of a spanoid \mathcal{S} is denoted by $\mathcal{C}_{\mathcal{S}}$ which is an intersection-closed family. Thus the family of all open sets which are complements of closed sets is a union-closed family and is denoted by $\mathcal{O}_{\mathcal{S}}$. One can construct all the derivations of the spanoid given its family of open or closed sets. Conversely, given any union-closed family of subsets of $[n]$, one can define a spanoid whose open sets are precisely the given family. Thus spanoids on $[n]$ are equivalent to union-closed families of subsets of $[n]$. The rank of a spanoid has a very simple interpretation in terms of its open sets, $\text{rank}(\mathcal{S})$ is equal to the size of the smallest hitting set for its family of open sets $\mathcal{O}_{\mathcal{S}}$. Moreover, $\text{rank}(\mathcal{S})$ is at most $\log |\mathcal{O}_{\mathcal{S}}|$. These connections are discussed in section 2.1.

Union-closed families are interesting combinatorial objects with a rich structure. The widely open Frankl's union-closed conjecture states that in every union-closed family of N sets, there exists an element which is contained in at least $N/2$ sets. Though this was proved for various special classes (see survey [BS15]), the best general bound is $\Omega(N/\log N)$ due to [Kni94, Wój99]. When seen in the framework of spanoids, this follows immediately from Claim 2.5 which says that there is a $\log N$ sized hitting set for every union-closed family of N sets. Thus there is an element which should hit at least $N/\log_2(N)$ sets. We hope that viewing union-closed families as spanoids could be of use in understanding them.

1.4. Organization. In section 2, we will present two alternative ways to represent spanoids that will be very useful. In section 3, we show upper bounds on the rank of q -LCSs for $q \geq 2$ thus proving the upper bounds in Theorems 3.5 and 3.7. In section 4, we construct q -LCSs thus proving the lower bound in Theorem 3.7. In section 5, we explore the relation between rank and f-rank of a spanoid. Specifically, we define the linear programs for $\text{LP}^{\text{entropy}}$ and LP^{cover} which provide upper and lower bounds on the f-rank , respectively. We will also calculate the rank and f-rank of the Pentagon spanoid defined in the introduction (Figure 1) and thus proving that f-rank can be strictly smaller. In section 6, we study products of spanoids which we use to amplify the gap between f-rank and rank from constant to polynomial.

2. Preliminaries on spanoids. We will now describe two equivalent ways to define spanoids which will turn out to be very useful.

2.1. Spanoids as union-closed or intersection-closed families. In this subsection, we will define an equivalent and more canonical way of describing spanoids in terms of intersection closed or union closed families. We have defined spanoids in Definition 1.5 by specifying the initial set of derivation rules such as $A \rightarrow i$. But two different initial sets of rules can lead to the same set of derivations and we should consider two spanoids to be equivalent if they lead to the same set of derivations. We will present an alternative way to describe spanoids which makes them equivalent to union-closed families of sets or alternatively intersection-closed families of sets. Moreover, this new representation is a more canonical way to represent spanoids since it will be based only on the set of derivations. For this, the main new notions we need

to define are that of a “closed set” and an “open set.”

DEFINITION 2.1 (closed and open sets). *Let \mathcal{S} be a spanoid on $[n]$. A closed set¹⁰ is a subset $B \subset [n]$ for which $\text{span}(B) = B$. A subset $B \subset [n]$ is called an open set if its complement is a closed set. The family of all closed sets of \mathcal{S} is denoted by $\mathcal{C}_{\mathcal{S}}$ and the family of all open sets of \mathcal{S} by $\mathcal{O}_{\mathcal{S}}$ (when it is clear from the context, we will drop the subscript \mathcal{S}).*

CLAIM 2.2. *In any spanoid \mathcal{S} on $[n]$,*

1. *the intersection of any number of closed sets is a closed set, i.e., $\mathcal{C}_{\mathcal{S}}$ is an intersection-closed family,*
2. *the union of any number of open sets is an open set, i.e., $\mathcal{O}_{\mathcal{S}}$ is a union-closed family and*
3. *for any set $A \subset [n]$, $\text{span}(A)$ is equal to the intersection of all closed sets containing A , i.e., $\text{span}(A) = \bigcap_{B \supset A, B \in \mathcal{C}_{\mathcal{S}}} B$.*

Proof. (1) Let $F = F_1 \cap F_2$ be the intersection of two closed sets. Suppose in contradiction that F spans some element $x \in [n] \setminus F$, then, by monotonicity, both F_1 and F_2 have to span x . Hence, $x \in \text{span}(F_1) \cap \text{span}(F_2) = F_1 \cap F_2 = F$ in contradiction.

(2) This just follows from (1) by taking complements.

(3) Let $F(A)$ be the intersection of all closed sets containing A . Since $\text{span}(A)$ is a closed set we clearly have $F(A) \subset \text{span}(A)$. To see the other direction, suppose $x \in \text{span}(A)$, and let F be any closed set containing A . Then, by monotonicity, F must also span x and so we must have $x \in F$. □

CLAIM 2.3. *A spanoid is uniquely determined by the set of all its closed (open) sets which is an intersection-closed (union-closed) family of subsets of $[n]$. Conversely, every intersection-closed (union-closed) family of subsets of $[n]$ defines a spanoid whose closed (open) sets are the given family.*

Proof. Given a spanoid \mathcal{S} on $[n]$, by Claim 2.2, we can define $\text{span}(A)$ in \mathcal{S} using just the closed sets as

$$\text{span}(A) = \bigcap_{B \supset A, B \in \mathcal{C}_{\mathcal{S}}} B.$$

And $A \models i$ in \mathcal{S} iff $i \in \text{span}(A)$. Thus given the set of all closed sets, we can reconstruct all the derivations of the spanoid.

For the converse, suppose we are given an intersection-closed family of subsets of $[n]$, say \mathcal{C} . We can define $\text{span}_{\mathcal{C}}(A) = \bigcap_{B \supset A, B \in \mathcal{C}} B$ and define a spanoid $\mathcal{S}_{\mathcal{C}}$ where $A \models i$ iff $i \in \text{span}_{\mathcal{C}}(A)$. It is easy to see that the closed sets of this spanoid $\mathcal{S}_{\mathcal{C}}$ is exactly \mathcal{C} . □

Thus an equivalent way to define a spanoid is to define all its closed (open) sets which is some intersection (union) closed family. The following claim shows that the rank of a spanoid has a very natural interpretation in terms of the open sets.

CLAIM 2.4. *The rank of a spanoid \mathcal{S} is the size of the smallest hitting set for the collection $\mathcal{O}_{\mathcal{S}} \setminus \{\emptyset\}$, i.e., a set which intersects every nonempty open set in $\mathcal{O}_{\mathcal{S}}$ nontrivially.*

Proof. Observe that a subset $A \subset [n]$ spans $[n]$ iff it is a hitting set for all the nonempty open sets in $\mathcal{O}_{\mathcal{S}}$. This is because if A doesn't hit some non-empty open set B , then A lies in the complement of B , i.e., $A \subset \bar{B}$. Since \bar{B} is closed, $\text{span}(A) \subset \bar{B} \neq [n]$. Conversely, if A hits every nonempty open set, then A should intersect

¹⁰Closed sets are analogous to “flats” or “subspaces” in matroids.

$[n] \setminus \text{span}(A)$ nontrivially or $\text{span}(A) = [n]$. Since $A \subset \text{span}(A)$, we should have $\text{span}(A) = [n]$. Therefore, $\text{rank}(\mathcal{S})$ is the size of the smallest hitting set for $\mathcal{C}_{\mathcal{S}}$. \square

This interpretation of the rank is used in section 5.2 to give an LP relaxation LP^{cover} which lower bounds the rank. We can also upper bound the rank of a spanoid in terms of the number of closed or open sets as the following claim shows.

CLAIM 2.5. *Let \mathcal{S} be a spanoid, then $\text{rank}(\mathcal{S}) \leq \log_2(|\mathcal{C}_{\mathcal{S}}|) = \log_2(|\mathcal{O}_{\mathcal{S}}|)$.*

Proof. Let $r = \text{rank}(\mathcal{S})$ and $R \subset [n]$ be a set of size $|R| = r$ spanning $[n]$. Since the rank of \mathcal{S} is r we know that R is independent (not spanned by any proper subset). For each of the 2^r subsets $S \in 2^R$ we consider the closed set $F_S = \text{span}(S)$. We claim that all of these are distinct. Suppose in contradiction that there were two distinct sets $S \neq T \in 2^R$ with $\text{span}(S) = \text{span}(T)$. Without loss of generality (w.l.o.g.) suppose there is an element $x \in T \setminus S$. Then $x \in \text{span}(S)$ and so we get that $R \setminus \{x\}$ spans R (by monotonicity) and so spans the entire spanoid in contradiction. Thus $|\mathcal{C}_{\mathcal{S}}| \geq 2^r$. \square

2.2. Spanoids as set systems. In this subsection, we will show yet another way of representing spanoids by families of sets. This representation (which is equivalent to spanoids) will be easier to work with and, in fact, we will later work almost exclusively with it instead of with the definition given in the introduction. Recall the notation introduced at the end of the introduction that, for sets S_1, \dots, S_n and for a subset $A \subset [n]$ we let $\cap A = \cap_{i \in A} S_i$.

DEFINITION 2.6 (intersection dimension of a set system). *The intersection-dimension of a family of sets S_1, \dots, S_n , denoted $\text{idim}(S_1, \dots, S_n)$ is the smallest integer d such that there exist a set $A \subset [n]$ of size d such that $\cap A = \cap [n]$.*

LEMMA 2.7 (set-representation of spanoids). *Let \mathcal{S} be a spanoid on $[n]$ with $\text{rank}(\mathcal{S}) = r$. Then there exists a family of sets S_1, \dots, S_n such that $A \models i$ in \mathcal{S} iff $\cap A \subset S_i$. In this case we say that the set family (S_1, S_2, \dots, S_n) is a set-representation of \mathcal{S} and this implies in particular that $\text{idim}(S_1, \dots, S_n) = \text{rank}(\mathcal{S})$.*

Proof. For $i \in [n]$ we define $S_i \subset \mathcal{C}_{\mathcal{S}}$ to be the subfamily of closed sets of \mathcal{S} containing the element $i \in [n]$. For the first direction of the proof suppose that A spans x in the spanoid \mathcal{S} . Then, by Claim 2.2, x belongs to any closed set containing A and so $\cap_{i \in A} S_i \subset S_x$. For the other direction, suppose $\cap_{i \in A} S_i \subset S_x$ or that any closed set containing A must also contain x . Hence, x is in the intersection of all closed sets containing A and, by Claim 2.2 we have that $x \in \text{span}(A)$. \square

An alternative way to represent spanoids is by unions. (T_1, T_2, \dots, T_n) is called a *union set-representation* of the spanoid \mathcal{S} when, $A \models i$ in \mathcal{S} iff $T_i \subset \cup_{j \in A} T_j$. Note that if (S_1, S_2, \dots, S_n) is an (intersection) set-representation for \mathcal{S} as in Lemma 2.7, then by taking complements, $(\bar{S}_1, \bar{S}_2, \dots, \bar{S}_n)$ is a union set-representation for \mathcal{S} and vice versa. Thus these two notions of representing a spanoid by sets are equivalent.

CLAIM 2.8. *Given a spanoid \mathcal{S} explicitly as a family of pairs $T \rightarrow i$ and some positive integer k , deciding if the rank of the spanoid is at most k is NP-complete.*

Proof. If we are given some subset A of size k along with derivations $A \models i$ for every element i of the spanoid, then we can check using basic rules that all of the derivations are correct. Note that the derivations are of polynomial length. Therefore, this problem is in NP.

To prove that it is NP-complete, we reduce set cover problem to this. Given a collection of sets $S_1, S_2, \dots, S_n \subset U$ such that $\cup_i S_i = U$ and some positive integer k , the set cover problem asks if there are at most k sets in the collection whose union is

U . To reduce it to the spanoid rank problem, we can create a spanoid \mathcal{S} over $[n] \sqcup U$ where the basic inference rules are given by the following:¹¹

1. $\{i\} \rightarrow a$ for all $a \in S_i$ and $i \in [n]$,
2. $S_i \rightarrow i$ for all $i \in [n]$.

Note that this spanoid has a polynomial length description in terms of the basic rules. We claim that the rank of this spanoid is equal to the smallest number of sets in the collection S_1, S_2, \dots, S_n which cover all of U . Let B be a subset of $[n] \sqcup U$ of size $\text{rank}(\mathcal{S})$ which spans all of $[n] \sqcup U$. Without loss of generality, we can assume that $B \subset [n]$. This is because if $a \in B$ for some $a \in U$, we can swap $a \in B$ with some $i \in [n]$ such that $S_i \ni a$. Now $\cup_{i \in B} S_i = U$ because $\text{span}(B) = [n] \sqcup U$. Conversely, if $\cup_{i \in B} S_i = U$ for some subset $B \subset [n]$, then $\text{span}(B) = U \sqcup [n]$. This completes the reduction. \square

3. Upper bounds on the rank of q -LCSs. In this section we prove the upper bounds on the rank of q -LCSs stated in Theorems 1.9 and 1.10. The proofs will rely on the set representation described in section 2.2 and on random restriction and contraction arguments given below.

3.1. Graph theoretic lemmas. In this subsection, we will prove a key technical lemma about a random graph process that will be useful for proving upper bounds on the rank of q -LCSs. We denote by $\mathcal{D}(n)$ the set of simple directed graphs on n vertices. We always assume w.l.o.g. that the set of vertices are the integers between 1 and n .

DEFINITION 3.1 ((α, β) -spread distribution). *Let μ be a distribution on $\mathcal{D}(n)$. We say that μ is (α, β) -spread if the following conditions are true for a graph G sampled from μ :*

1. *Each vertex $i \in [n]$ has an incoming edge with probability at least α , i.e.,*

$$\forall i \Pr_{G \sim \mu} [\exists j : (j, i) \in E(G)] \geq \alpha.$$

2. *For every $i, j \in [n]$, the probability that (j, i) is an edge is at most β/n , i.e.,*

$$\forall i, j \Pr_{G \sim \mu} [(j, i) \in E(G)] \leq \frac{\beta}{n}.$$

For example, one can generate an $(k/n, 1)$ -spread distribution μ on $\mathcal{D}(n)$ in the following way: Fix arbitrary sets $S_1, \dots, S_n \subset [n]$ of size k each. To sample a graph G from μ , pick a uniformly random element $j \in [n]$, and let G be the directed graph containing the edges (j, i) for each i such that $j \in S_i$. This satisfies the definition since for any fixed $i \in [n]$, i has an incoming edge if $j \in S_i$ which happens with probability $|S_i|/n = k/n$. And for any fixed $i', j' \in [n]$, the probability that (j', i') is an edge is at most $1/n$ since this happens only when $j' = j$ and j is chosen uniformly at random from $[n]$. Note that the sampled edges overall are highly correlated (they all have j as an endpoint).

We will need following simple observation about (α, β) -spread distributions.

LEMMA 3.2. *Let μ be an (α, β) -spread distribution on $\mathcal{D}(n)$. For every vertex i and every subset $S \subset [n]$ of size at most $\frac{\alpha n}{2\beta}$,*

$$\Pr_{G \sim \mu} [\exists j \notin S : (j, i) \in E(G)] \geq \frac{\alpha}{2}.$$

¹¹ U and $[n]$ are considered disjoint.

Proof. The proof follows from union bound and properties of (α, β) -spread distributions:

$$\begin{aligned} \alpha &\leq \Pr_{G \sim \mu}[\exists j : (j, i) \in E(G)] \\ &\leq \Pr[\exists j \in S : (j, i) \in E(G)] + \Pr[\exists j \notin S : (j, i) \in E(G)] \\ &\leq \sum_{j \in S} \Pr[(j, i) \in E(G)] + \Pr[\exists j \notin S : (j, i) \in E(G)] \\ &\leq \frac{\alpha n}{2\beta} \cdot \frac{\beta}{n} + \Pr[\exists j \notin S : (j, i) \in E(G)] \\ &= \frac{\alpha}{2} + \Pr[\exists j \notin S : (j, i) \in E(G)]. \quad \square \end{aligned}$$

Given a distribution μ on graphs we would like to study the random process in which we, at each iteration, sample from μ and “add” the edges we got to the graph obtained so far. For two graphs G and H on the same set of vertices, we denote by $G \cup H$ their set theoretic union (as a union of edges).

DEFINITION 3.3 (graph process associated with μ). *Let μ be a distribution on $\mathcal{D}(n)$. We define a sequence of random variables G_t^μ , $t = 0, 1, 2, \dots$, as follows: G_0^μ is the empty graph on $[n]$ vertices. At each step $t \geq 1$ we sample a graph G according to μ (independently from all previous samples) and set $G_t = G_{t-1} \cup G$.*

For a graph $G \in \mathcal{D}(n)$ and a vertex $i \in [n]$ we denote by $\text{Rea}(i)$ the set of vertices that are reachable from i (via walking on directed edges). By convention, a vertex is always reachable from itself. Similarly, for a set of vertices $S \subset [n]$ we denote by $\text{Rea}(S) = \cup_{i \in S} \text{Rea}(i)$ the set of vertices reachable from some vertex in S . We denote the set of strongly connected components of G by $\Gamma(G)$. We denote by $C(i) \in \Gamma(G)$ the strongly connected component of G containing i . We say that $C \in \Gamma(G)$ is a *source* if C has no incoming edges from any vertex not in C .

LEMMA 3.4. *Let μ be an (α, β) -spread distribution on $\mathcal{D}(n)$, and let G_t^μ be its associated graph process. Then, for all $t \geq 0$, there is positive probability that the graph $\Gamma(G_t^\mu)$ has at most*

$$n \cdot (1 - \alpha/4)^t + \frac{2\beta}{\alpha}$$

sources.

Proof. If $C \in \Gamma(G)$ is a source, we define the *weight* of C to be the number of vertices reachable from C (including vertices of C) that are not reachable from any other source of G . More formally, let

$$\text{Rea}'(C) = \{j \in \text{Rea}(C) \mid j \notin \text{Rea}(C') \ \forall \text{sources } C' \in \Gamma(G), C' \neq C\}.$$

Then the weight of a source $C \in \Gamma(G)$ is denoted by $w(C) = |\text{Rea}'(C)|$ (we do not define weight for components that are not sources). Let us call a source $C \in \Gamma(G_t)$ “light” if its weight $w(C)$ is at most $k = \frac{\alpha n}{2\beta}$ and “heavy” otherwise. By the definition of weight, there could be at most $n/k = 2\beta/\alpha$ heavy sources.

We will argue that, in each step, as we move from G_t to G_{t+1} , the number of light sources must decrease by a factor of $(1 - \alpha/4)$ with positive probability. For that purpose, suppose there are m_t sources in G_t and among them m'_t are light. Fix some light source and pick a representative vertex i from it. Since i is contained in a light source, $|\text{Rea}'(C(i))| \leq \frac{\alpha n}{2\beta}$. When going to $G_{t+1} = G_t \cup G$, i gets an incoming edge from outside the set $\text{Rea}'(C(i))$ with probability at least $\alpha/2$ by Lemma 3.2. If

this happens then in G_{t+1} , this source will either stop being a source or merge with another source.

Picking a representative for each light source in G_t , we see that the expected number of representatives i which get a new incoming edge from outside $\text{Rea}'(C(i))$ is at least $(\alpha/2)m'_t$. Hence, this quantity is obtained with positive probability. Now, if at least $(\alpha/2)m'_t$ light sources “merge” with another source or stop being a source in G_{t+1} , then the total number of light sources must decrease by at least $(\alpha/4)m'_t$ (the worst case being that $(\alpha/4)m'_t$ disjoint pairs of light sources merge with each other). Hence, with positive probability we get that $m'_{t+1} \leq m'_t \cdot (1 - \alpha/4)$. Therefore, since the samples in each step t are independent, there is also a positive probability that $m'_t \leq n \cdot (1 - \alpha/4)^t$ and $m_t \leq m'_t + 2\beta/\alpha$. This completes the proof. \square

3.2. Proof of upper bound from Theorem 1.9.

THEOREM 3.5 (rank of 2-LCSs). *Let \mathcal{S} be a 2-LCS on $[n]$ with error-tolerance δ . Then $\text{rank}(\mathcal{S}) \leq O(\frac{1}{\delta} \log_2 n)$.*

Proof. We will work with the (equivalent) set formulation: Let $\mathcal{F} = \{S_1, \dots, S_n\}$ be a set system representing the spanoid \mathcal{S} as in Lemma 2.7.

We start by defining an (α, β) -spread distribution μ on $\mathcal{D}(n)$ as follows: To sample a graph G from μ we first pick $\ell \in [n]$ uniformly at random. Then we add a directed edge from j to i for every i, j such that $\{j, \ell\} \in M_i$. In this case we have $S_j \cap S_\ell \subseteq S_i$ and so, after restricting to S_ℓ we have $S_j \cap S_\ell \subseteq S_i \cap S_\ell$.

CLAIM 3.6. *μ is a $(2\delta, 1)$ -spread distribution.*

Proof. For any fixed $i \in [n]$, i will get an incoming edge if ℓ , which is randomly chosen from $[n]$, belongs to M_i . Since M_i has at least δn edges, this will happen with probability at least 2δ . Now fix any $i, j \in [n]$, (j, i) will be an edge iff ℓ is equal to the vertex that matches j in the matching M_i , this happens with probability at most $1/n$. If j is not matched in M_i , the probability is zero. \square

Consider the graph process G_t^μ and let $S_{\ell_1}, \dots, S_{\ell_t}$ be the sets chosen in the t iterations of sampling from μ . If $i \in \text{Rea}(j)$ in the graph G_t^μ , this means that, after restricting to the intersection $S = S_{\ell_1} \cap \dots \cap S_{\ell_t}$, the set S_j is contained in S_i (i.e., $S_j \cap S \subseteq S_i \cap S$). By Lemma 3.4, after $t = O(\frac{1}{\delta} \log_2 n)$ steps, the graph process G_t^μ will contain $r = O(1/\delta)$ sources. Pick a representative S_{a_1}, \dots, S_{a_r} from each of these sources. Then, the intersection of the $t + r = O(\frac{1}{\delta} \log_2 n)$ sets $S_{\ell_1}, \dots, S_{\ell_t}$ and S_{a_1}, \dots, S_{a_r} is contained in all n sets S_1, \dots, S_n . That is because, when restricted to the intersection of $S_{\ell_1}, \dots, S_{\ell_t}$, each set S_i contains one of the sets $S_{a_j}, j \in [r]$. \square

3.3. Proof of upper bound from Theorem 1.10.

THEOREM 3.7 (rank of q -query LCSs). *Let \mathcal{S} be a q -LCS with error-tolerance δ and $q \geq 3$. Then*

$$\text{rank}(\mathcal{S}) \leq O\left(\delta^{-\frac{1}{q-1}} \cdot n^{\frac{q-2}{q-1}} \log_2 n\right).$$

Proof. Like the 2-query case, we work with the set representation $\mathcal{F} = \{S_1, \dots, S_n\}$ of \mathcal{S} as in Lemma 2.7. We follow the same strategy as in the proof of the 2-query case. The difference is that, in this case, we will need to pick many sets to restrict to in each step instead of just one. The first observation is that, if $\{j_1, \dots, j_q\} \in M_i$, then, restricted to the intersection $S = S_{j_1} \cap \dots \cap S_{j_{q-1}}$ we have $S_{j_q} \subseteq S_i$. The second observation is that, if we choose a subset $J \subset [n]$ of size roughly $n^{\frac{q-2}{q-1}}$, then, in expectation, J will contain $q - 1$ elements in one of the q -subsets of M_i for a constant fraction of

the i 's. Repeating this a logarithmic number of times and using Lemma 3.4, as in the proof of Theorem 3.5 will then complete the proof.

We start by defining an (α, β) -spread distribution μ on $\mathcal{D}(n)$. To sample a graph G from μ we first pick a random set $J \subset [n]$ such that each $j \in [n]$ is chosen to be in J independently with probability $(\delta n)^{-1/(q-1)}$. By Markov's inequality we have that

$$(3.1) \quad \Pr \left[|J| \geq 4 \cdot \delta^{-\frac{1}{q-1}} n^{\frac{q-2}{q-1}} \right] \leq 1/4.$$

For each $i \in [n]$ and each q -subset $T \in M_i$ we select $q-1$ elements of T arbitrarily and refer to them as the *distinguished $(q-1)$ -subset of T* . We now argue that, for each $i \in [n]$, there is relatively high probability that J will contain the distinguished $(q-1)$ -subset of at least one q -subset in M_i .

CLAIM 3.8. *Let E_i denote the event that J contains the distinguished $(q-1)$ -subset from at least one q -subset in M_i . Then, for each $i \in [n]$ we have that $\Pr[E_i] \geq 1/2$.*

Proof. J will contain the distinguished $q-1$ elements in a specific q -subset with probability $(\delta n)^{-1}$. Since the δn q -subsets in M_i are disjoint, the probability that J will not contain any of the distinguished $(q-1)$ -subsets is at most $(1 - (1/\delta n))^{\delta n} \leq 1/2$. \square

We are now ready to define the edges in the graph G sampled by μ . First, we check if $|J| \geq 4 \cdot \delta^{-\frac{1}{q-1}} n^{\frac{q-2}{q-1}}$. If this is the case, then μ outputs the empty graph (by (3.1) this happens with probability at most $1/4$). Otherwise, for each $i \in [n]$ we check to see if J contains the distinguished $(q-1)$ -subset from one of the q -subsets of M_i . If there is at least one such q -subset, we pick one of them uniformly at random. Suppose the q -subset we chose is $\{j_1, \dots, j_q\}$ and that the distinguished elements are the first $q-1$. Then we add the directed edge $j_q \rightarrow i$ to the graph G . By the above discussion, we know that, restricted to the intersection of all sets indexed by J the set S_{j_q} is contained in S_i (hence the directed edge representing set inclusion).

CLAIM 3.9. μ is $(1/4, 1/\delta)$ -spread.

Proof. By Claim 3.8, and since the probability that $|J|$ is too large is at most $1/4$ we see that any fixed $i \in [n]$ will get an incoming edge in G with probability at least $\alpha = 1/4$. For any fixed $i, j \in [n]$, since the distribution of the special q -subset, which is contributing an edge to i is uniform in M_i (conditioned on J containing a q -subset from M_i), we can conclude that $(j, i) \in E(G)$ with probability at most $1/(\delta n) = \beta/n$. This proves the claim. \square

Now, applying Lemma 3.4, we get that, after $t = O(\log_2 n)$ steps, the graph process G_t^μ will contain at most $O(1/\delta)$ sources with positive probability. Let J_1, \dots, J_t be the sets chosen in the different steps of the process and, w.l.o.g., remove any of them that were too big (i.e., when the graph sampled by μ was empty). Hence, all of the sets satisfy $|J_i| \leq 4 \cdot \delta^{-\frac{1}{q-1}} n^{\frac{q-2}{q-1}}$. Now, let S be the intersection of all sets S_j such that j belongs to at least one of the sets J_i . Then, restricted to S , each of the sets S_i contains one of the sources in the graph G_t^μ . Hence, if we add to our intersection a representative from each of the sources, we will get a set that is contained in all of the sets S_j . The total number of sets we end up intersecting is bounded by

$$O(1/\delta) + \sum_{i=1}^t |J_i| = O\left(\delta^{-\frac{1}{q-1}} \cdot n^{\frac{q-2}{q-1}} \log_2 n\right).$$

This completes the proof of the theorem. \square

4. Constructing q -LCSs with high rank. In this section we prove the lower bound part of Theorem 1.10 (the lower bound for the 2-query case follows from the Hadamard code construction). We will in fact generate this spanoid at random by picking, for each $i \in [n]$, a random q -matching M_i on $[n]$ and, for each q -subset $T \in M_i$ add the rule $T \rightarrow i$. The resulting spanoid will thus have, by design, the structure of a q -LCS. The reason why this spanoid should have high rank (with high probability) relies on the following observation. Suppose $A \subset [n]$ is a set that spans $[n]$. This means that there is a sequence of derivations $T_i \rightarrow i$ with each q -subset T_i in the matching M_i that eventually generates all of $[n]$. We can limit ourselves to the first $C \cdot |A|$ such derivations for $C = \Theta(\log n)$. These derivations generate a set A' of size $(C + 1)|A|$ (including the original A and the $C|A|$ newly derived elements). Now, the set A' must contain all of the q -subsets T_i for $C|A|$ values of i . However, the union of randomly chosen $C|A|$ q -subsets will generally have size much larger than $(C + 1)|A|$ (closer to $q \cdot C|A|$).

THEOREM 4.1 (existence of high rank q -LCSs). *For any integer $q \geq 3$ and all sufficiently large n the following holds. Consider the following distribution generating a spanoid \mathcal{S} on base set $[n]$. For each $i \in [n]$ pick a q -matching M_i of size $\lfloor n/q \rfloor$ uniformly at random and add the rule $T \rightarrow i$ for all $T \in M_i$. Then, with probability approaching one, $\text{rank}(\mathcal{S})$ is larger than $r = cn^{\frac{q-2}{q-1}} / \log_2(n)$, where $0 < c < 1$ is an absolute constant.*

Proof. Let $m = r \cdot \log_2(n) = cn^{\frac{q-2}{q-1}}$. If the rank of \mathcal{S} is at most r , then there exists a set $A \subset [n]$ of size r that spans (using the rules obtained from the n random matchings M_1, \dots, M_n) the entire base set $[n]$. We will upper bound the probability that such a set exists by bounding the bigger event given by the existence of a set of m rules that can be applied one after another starting with the original set A . That is, let \mathcal{E} denote the event that there exists a set A of size r on which one can sequentially apply m rules of the form $T_{j_i} \rightarrow j_i$ with each T_{j_i} belonging to the matching M_{j_i} and for m different values $j_1, \dots, j_m \in [n]$ arriving at the final set $\hat{A} = A \cup \{j_1, \dots, j_m\}$. If A spans $[n]$, then clearly the event \mathcal{E} must hold and so, it is enough to show that \mathcal{E} happens with probability approaching zero.

We will present the event \mathcal{E} as the union of (possibly overlapping) smaller events and then use the union bound, bounding the probability that each one occurs and multiplying by the number of bad events. For a matching M_i (with some fixed ordering on its edges) and $\ell \in [n/q]$, let $M_j(\ell)$ be the ℓ th edge in the matching M_j . Note that the marginal distribution of $M_j(\ell)$ for a fixed j, ℓ is uniformly distributed among all q -subsets of $[n]$. Given a set $A \subset [n]$ of size r , a subset of m indices $J = \{j_1, j_2, \dots, j_m : j_1 < j_2 < \dots < j_m\} \subset [n]$ and a sequence $L = (\ell_1, \ell_2, \dots, \ell_m) \in ([n/q])^m$, denote by $\mathcal{E}(A, J, L)$ the event in which the set A spans the set $\hat{A} = A \cup J$ using the rules $M_{j_i}(\ell_i) \rightarrow j_i$ applied in *some* order. In particular this implies that $M_{j_i}(\ell_i) \subset \hat{A}$ for $i \in [m]$. Therefore, for every fixing of A, J, L , we can bound

$$\Pr[\mathcal{E}(A, J, L)] \leq \Pr[\forall i \in [m] M_{j_i}(\ell_i) \subset \hat{A}] = \prod_{i=1}^m \Pr[M_{j_i}(\ell_i) \subset \hat{A}],$$

where we used the independence of M_1, M_2, \dots, M_n . Since the marginal distribution of $M_{j_i}(\ell_i)$ is uniform over all q -sized subsets,

$$\Pr[M_{j_i}(\ell_i) \subset \hat{A}] \leq \frac{\binom{m+r}{q}}{\binom{n}{q}} \leq \left(\frac{4m}{n}\right)^q.$$

Taking the product over all $i \in [m]$, we get

$$\Pr[\mathcal{E}(A, J, L)] \leq \left(\frac{4m}{n}\right)^{qm}.$$

To complete the proof we can upper bound the number of tuples (A, J, L) as above by

$$\binom{n}{r} \cdot \binom{n}{m} \cdot \left[\frac{n}{q}\right]^m \leq n^r \cdot \left(\frac{\epsilon n}{m}\right)^m \cdot n^m \leq \left(\frac{6n^2}{m}\right)^m,$$

where the last inequality used the fact that $r/m \leq 1/\log_2(n)$. Putting these bounds together we get that

$$\Pr[\mathcal{E}] \leq \left(\frac{4m}{n}\right)^{qm} \left(\frac{6n^2}{m}\right)^m = \left(\frac{6 \cdot 4^q \cdot m^{q-1}}{n^{q-2}}\right)^m$$

which is exponentially decreasing in m for the given choice of $m = c \cdot n^{(q-2)/(q-1)}$ and for c a sufficiently small constant. \square

Explicit q -LCSs. One could ask for a more explicit construction of a q -LCS with rank equal to (or even close to) that stated above. We are not able to give such a construction but can relate this problem to a longstanding open problem in explicit construction of expander graphs. A bipartite (balanced) expander of degree q is a bipartite graph with n left vertices L and n right vertices R such that the degree of each vertex is q and such that sets $A \subset L$ of size “not too large” have many neighbors in R . More specifically, one typically asks that sets with $|A| \leq n/2$ have at least $(1 + \epsilon)|A|$ right neighbors for some constant $\epsilon > 0$. It is quite easy to see that a random graph of this form will be a good expander with high probability and, by now, there are also many explicit constructions [HLW06]. One can also consider *unbalanced* bipartite expanders in which $|L| \gg |R|$. Take, for example, the setting in which $|L| = n^2, |R| = n$, and when the degree of every vertex in L is some constant q . A simple probabilistic argument shows that sufficiently small sets in L , namely sets of size $|A| \leq n^{\alpha_q}$ with $\alpha_q < 1$ a constant depending on q and approaching 1 as q grows, have many neighbors in R (say, at least $2|A|$). However, no explicit constructions of such graphs are known (for any constant q and any $\alpha_q > 0$). The property we needed in our random construction of LCSs can be thought of as an “easier” variant of the expander construction problem. Given q -matchings M_1, \dots, M_n each of size δn consider the bipartite graph with $L = [n] \times [\delta n]$ and $R = [n]$. We identify each vertex $(i, j) \in L$ with the j th q -subset T_{ij} of M_i and connect it to the q neighbors in R given by that q -subset. For our proof to work we need the property that there is no small set containing many q -subsets from different matchings. This corresponds to asking for the above graph to be an expander for a restricted family of sets, namely to sets that have at most one vertex (i, j) for a given i (with each subgraph $(i, *)$ defining a matching).

5. Functional rank versus spanoid rank. In this section we analyze the five element spanoid Π_5 described in the introduction (Figure 1) and show that its rank is strictly larger than its functional rank. Along the way we formulate the LP relaxation $\text{LP}^{\text{cover}}(\mathcal{S})$ which lower bounds the functional rank in general and another linear program $\text{LP}^{\text{entropy}}(\mathcal{S})$ which upper bounds the functional rank.

First, we give the lower bound by constructing a consistent code over an alphabet of size 4 with $32 = 4^{2.5}$ codewords.

CLAIM 5.1. *Let Π_5 be the pentagon spanoid defined in Figure 1. Then, $f\text{-rank}(\mathcal{M}_5) \geq 2.5$.*

Proof. We will construct a consistent code over the alphabet $\Sigma = \{0, 1\}^2$. Each codeword will be indexed by an element of $\{0, 1\}^5$. The codeword $w(x) \in \Sigma^5$ corresponding to $(x_1, x_2, x_3, x_4, x_5) \in \{0, 1\}^5$ will be $((x_5, x_2), (x_1, x_3), (x_2, x_4), (x_3, x_5), (x_4, x_1))$ as shown in Figure 2. In other words, we place the bits x_1, \dots, x_5 on the vertices of the cycle and then assign to each vertex the symbol of $\Sigma = \{0, 1\}^2$ comprised of the bits of its two neighbors on the cycle. It is now straight forward to verify that one can compute the coordinate $w_i, i \in [5]$ from the two coordinates spanning it in Π_5 . For example, the span rule $\{1, 2\} \rightarrow 4$ requires us to compute $w_4 = (x_3, x_5)$ from $w_1 = (x_5, x_2)$ and $w_2 = (x_1, x_3)$, which can be easily done (by symmetry, this is the situation in all of the other rules). \square




FIG. 2. A consistent code for Π_5 ; each coordinate can be recovered from the coordinates of the opposite edge.

5.1. An upper bound on functional rank via $\text{LP}^{\text{entropy}}$. In this section, we will give an LP upper bound for $f\text{-rank}(\mathcal{S})$ using properties of Shannon entropy. We will show that this upper bound matches the lower bound of 2.5 for $f\text{-rank}(\Pi_5)$ shown in Claim 5.1, thus proving that $f\text{-rank}(\Pi_5) = 2.5$. We will begin by recollecting some properties of Shannon entropy.

Given a random variable X supported on some domain A , its (Shannon) entropy is defined as

$$H(X) = - \sum_{a \in A} \Pr[X = a] \log(\Pr[X = a]).$$

The Shannon entropy of a random variable measures its information content. The conditional entropy of X given an other random variable Y is defined as $H(X|Y) = H(X, Y) - H(Y)$. And the conditional mutual information between X and Y given a third random variable Z is defined as $I(X : Y|Z) = H(X|Z) + H(Y|Z) - H(XY|Z)$. Equivalently, $I(X : Y|Z) = H(X, Z) + H(Y, Z) - H(X, Y, Z) - H(Z)$. Shannon proved that conditional entropy and conditional mutual information are always non-negative [CT91]. These are called basic information inequalities.

Let $X = (X_1, X_2, \dots, X_n)$ be a random variable made up of n coordinates. Define a function $f : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ as $f(S) = H(X_S)$, where $X_S = (X_i)_{i \in S}$. Note that f is a *monotone increasing* function because

$$f(A \cup B) - f(A) = H(X_{A \cup B}) - H(X_A) = H(X_A, X_B) - H(X_B) = H(X_A|X_B) \geq 0.$$

Moreover, f is a *submodular function*, i.e., for every $A, B \subset [n]$, $f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$. This is because

$$\begin{aligned} 0 &\leq I(X_{A \setminus B} : X_{B \setminus A} | X_{A \cap B}) \\ &= H(X_{A \setminus B}, X_{A \cap B}) + H(X_{B \setminus A}, X_{A \cap B}) - H(X_{A \setminus B}, X_{B \setminus A}, X_{A \cap B}) - H(X_{A \cap B}) \\ &= H(X_A) + H(X_B) - H(X_{A \cup B}) - H(X_{A \cap B}) \\ &= f(A) + f(B) - f(A \cup B) - f(A \cap B). \end{aligned}$$

In fact, the monotone increasing submodular property of f captures all of the inequalities that can be obtained by using the basic information inequalities. But when $n \geq 4$, the entropies $H(X_S)$ satisfy some extra linear inequalities that are not captured by the basic information inequalities. These mysterious inequalities are called non-Shannon type inequalities and a few such inequalities are known [ZY98], but they are not well understood. The set

$$\Gamma_n^* = \{(H(X_S))_{S \subset [n], S \neq \emptyset} : X = (X_1, X_2, \dots, X_n) \text{ r.v.}\},$$

where X ranges over all n jointly distributed random variables, is called the *entropic region* for n random variables. Γ_n^* is a convex cone, but neither Γ_n^* nor its closure $\overline{\Gamma_n^*}$ are polyhedral for $n \geq 4$ [Mat07], i.e., they are not defined by a finite number of linear inequalities. See [Yeu08] for more information about non-Shannon type information inequalities and the entropic region.

We are now ready to set up the linear program for upper bounding the functional rank of a spanoid. Let \mathcal{S} be a spanoid on $[n]$, and let $C \subset \Sigma^n$ be a code over some alphabet Σ which is consistent with the spanoid \mathcal{S} , i.e., whenever $A \rightarrow i$ in the spanoid, for every codeword $c \in C$, c_i is determined by $c|_S$. Let $X = (X_1, X_2, \dots, X_n)$ be a random variable with uniform distribution over C . Then $A \rightarrow i$ in \mathcal{S} implies that $H(X_{A \cup \{i\}}) = H(X_A)$. The dimension of the code C is

$$k = \frac{\log |C|}{\log |\Sigma|} = \frac{H(X)}{\log |\Sigma|}.$$

So upper bounding the functional rank of \mathcal{S} is equivalent to upper bounding $H(X_1, X_2, \dots, X_n)$ where X is a random variable distributed over Σ^n such that $H(X_{A \cup \{i\}}) = H(X_A)$ whenever $A \rightarrow i$ in \mathcal{S} . Define $f : 2^{[n]} \rightarrow \mathbb{R}$ by

$$f(S) = \frac{H(X_S)}{\log |\Sigma|}.$$

Clearly $f(\emptyset) = 0$ and $f(\{i\}) \leq 1$ for all $i \in [n]$. The basic information inequalities are equivalent to saying that f is a *monotone increasing submodular function*, i.e., for every subsets $A, B \subset [n]$,

$$f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$$

and if $A \subset B$, then $f(A) \leq f(B)$. Thus the best upper bound we can derive on $H(X)$ using the basic information inequalities is captured by the following linear program:

$$\begin{aligned}
 \text{LP}^{\text{entropy}}(\mathcal{S}) &= \max f([n]), \\
 f(\phi) &= 0, \\
 f(\{i\}) &\leq 1 \quad \forall i \in [n], \\
 f(A \cup B) + f(A \cap B) &\leq f(A) + f(B) \quad \forall A, B \subset [n], \\
 f(A) &\leq f(B) \quad \forall A \subset B \subset [n], \\
 f(A \cup \{i\}) &= f(A) \text{ whenever } A \models i \text{ in } \mathcal{S}.
 \end{aligned}
 \tag{5.1}$$

Note that $\text{LP}^{\text{entropy}}(\mathcal{S})$ is always at most $\text{rank}(\mathcal{S})$. This is because, if $A \subset [n]$ is such that $\text{span}(A) = [n]$, then any feasible f in the LP (5.1) should satisfy $f([n]) \leq |A|$. The following claim formally states that $\text{LP}^{\text{entropy}}(\mathcal{S})$ upper bounds the functional rank of the spanoid \mathcal{S} and lower bounds $\text{rank}(\mathcal{S})$, the proof of which follows immediately from the above discussion.

CLAIM 5.2. *For any spanoid \mathcal{S} , $\text{f-rank}(\mathcal{S}) \leq \text{LP}^{\text{entropy}}(\mathcal{S}) \leq \text{rank}(\mathcal{S})$.*

Note that the functional rank could be smaller than $\text{LP}^{\text{entropy}}(\mathcal{S})$ (though we do not know of an explicit example). This is because the basic information inequalities do not characterize the entropic region of more than three random variables. It might be possible to obtain better upper bounds on the functional rank by using non-Shannon type information inequalities. But in the case of the pentagon spanoid Π_5 defined in Figure 1, we will show that $\text{LP}^{\text{entropy}}(\mathcal{S})$ gives the tight upper bound.

CLAIM 5.3. *$\text{f-rank}(\Pi_5) = \text{LP}^{\text{entropy}}(\Pi_5) = 2.5 < 3 = \text{rank}(\Pi_5)$.*

Proof. We will show that value of the LP in (5.1) is at most 2.5. By the spanoid rules, we know that $f(\{i, i + 2, i + 3\}) = f(\{i + 2, i + 3\})$ for every $i \in [5]$ where the addition is modulo 5. We want to upper bound $f([5])$. By submodularity of f ,

$$\begin{aligned}
 f(\{1, 3, 4\}) + f(\{1, 2, 4\}) &\geq f(\{1, 4\}) + f(\{1, 2, 3, 4\}) \\
 f(\{1, 4\}) + f(\{1, 5\}) &\geq f(\{1\}) + f(\{1, 4, 5\}).
 \end{aligned}$$

By the inference rules, $f(\{1, 3, 4\}) = f(\{3, 4\})$, $f(\{1, 2, 4\}) = f(\{1, 2\})$ and $f(\{1, 2, 3, 4\}) = f(\{1, 4, 5\}) = f(\{1, 2, 3, 4, 5\})$. Therefore, the above two inequalities imply

$$f(\{1, 2\}) + f(\{1, 5\}) + f(\{3, 4\}) \geq f(\{1\}) + 2f(\{1, 2, 3, 4, 5\}).$$

By rotational symmetry, we can obtain five inequalities of this form. Summing them, and observing that every adjacent pair of vertices is counted three times on the left hand side, we get

$$3 \sum_i f(\{i, i + 1\}) \geq \sum_i f(\{i\}) + 10f(\{1, 2, 3, 4, 5\}).$$

Upper bounding $f(\{i, j\})$ by $f(\{i\}) + f(\{j\})$, we get

$$\begin{aligned}
 6 \sum_i f(\{i\}) &\geq \sum_i f(\{i\}) + 10f(\{1, 2, 3, 4, 5\}) \\
 \Rightarrow f(\{1, 2, 3, 4, 5\}) &\leq \frac{1}{2} \sum_i f(\{i\}) \leq \frac{5}{2}.
 \end{aligned}$$

Therefore, by Claim 5.2, $\text{f-rank}(\Pi_5) \leq \text{LP}^{\text{entropy}}(\Pi_5) \leq 2.5$. By Claim 5.1, $\text{f-rank}(\Pi_5) \geq 2.5$. This implies the required claim. \square

5.2. A lower bound on functional rank via LP^{cover} . In this section, we will prove lower bounds on functional rank by constructing consistent codes. And the best code one can construct in this way is captured by a very natural LP relaxation of the spanoid rank called LP^{cover} . The code constructed in Claim 5.1 for the pentagon spanoid Π_5 can be viewed as an instance of a more general scheme based on a union set-representation of a spanoid.

CONSTRUCTION 5.4. *Let (S_1, S_2, \dots, S_n) be a union set-representation for the spanoid \mathcal{S} where S_1, \dots, S_n are subsets of a universe $U = \cup_i S_i$ and each S_i is of size at most ℓ . So whenever $T \models i$ in \mathcal{S} , $S_i \subset \cup_{t \in T} S_t$. Such a representation can be used to define a consistent code C of dimension $|U|/\ell$ as follows. The codewords are images of the map $C : \{0, 1\}^U \rightarrow \Sigma^n$, where $\Sigma = \{0, 1\}^\ell$ given by $C(x)_i = (x_u)_{u \in S_i}$. It is easy to check that this indeed gives a code consistent with \mathcal{S} and the dimension of the code is $|U|/\ell$.*

We will now show that the best consistent code (i.e., of highest dimension) based on this approach can be characterized by an LP. Note that by Claim 2.4, the rank of a spanoid \mathcal{S} , is the size of the smallest hitting set for \mathcal{O} , the set of all open sets of \mathcal{S} . We can write an LP relaxation for the smallest hitting set for \mathcal{O} . Let $\mathcal{O}^* \subset \mathcal{O}$ be the set of minimal open sets, it is enough to hit every set in \mathcal{O}^* .

$$\begin{aligned}
 \text{LP}^{\text{cover}}(\mathcal{S}) &= \min \sum_{i=1}^n x_i, \\
 (5.2) \quad &x_i \geq 0, \\
 &\sum_{i \in S} x_i \geq 1 \quad \forall S \in \mathcal{O}^*.
 \end{aligned}$$

Since the LP is a relaxation,

$$\text{rank}(\mathcal{S}) \geq \text{LP}^{\text{cover}}(\mathcal{S}).$$

We can round an LP solution to get an integral hitting set by losing a factor of $O(\text{VC-dim}(\mathcal{O}^*) \cdot \log(\text{LP}^{\text{cover}}(\mathcal{S})))$ where $\text{VC-dim}(\mathcal{O}^*)$ is the VC-dimension of \mathcal{O}^* [ERS05, BG95]. Therefore,

$$1 \leq \frac{\text{rank}(\mathcal{S})}{\text{LP}^{\text{cover}}(\mathcal{S})} \lesssim \text{VC-dim}(\mathcal{O}^*) \cdot \log(\text{LP}^{\text{cover}}(\mathcal{S})).$$

Note that $\text{VC-dim}(\mathcal{O}^*) \leq \log |\mathcal{O}^*|$ always. By LP duality, we can write a dual LP for $\text{LP}^{\text{cover}}(\mathcal{S})$ with a dual variable λ_S for every $S \in \mathcal{O}^*$,

$$\begin{aligned}
 \text{LP}^{\text{cover}}(\mathcal{S}) &= \max \sum_{S \in \mathcal{O}^*} \lambda_S, \\
 (5.3) \quad &\lambda_S \geq 0, \\
 &\sum_{S \ni i} \lambda_S \leq 1 \quad \forall i \in [n].
 \end{aligned}$$

The following lemma shows that $\text{LP}^{\text{cover}}(\mathcal{S})$ is the largest dimension of a consistent code one can obtain using union set-representation of a spanoid as shown in Construction 5.4. Since the construction of a consistent code for the pentagon spanoid Π_5 in Claim 5.1 is obtained by Construction 5.4 and since we know that $\text{f-rank}(\Pi_5) = 2.5$, it turns out that $\text{LP}^{\text{cover}}(\Pi_5) = 2.5$.

LEMMA 5.5. *Let C be a code consistent with \mathcal{S} obtained by using Construction 5.4, then $\dim(C) \leq \text{LP}^{\text{cover}}(\mathcal{S})$. Moreover, there exists a code C obtained by using Construction 5.4 giving equality and so*

$$f\text{-rank}(\mathcal{S}) \geq \text{LP}^{\text{cover}}(\mathcal{S}).$$

Proof. We will first show that one can get a code whose rank is at least $\text{LP}^{\text{cover}}(\mathcal{S})$. Let \mathcal{S} be a spanoid on n elements with minimal open sets \mathcal{O}^* . Let $\lambda_{\mathcal{S}}$ be the optimal solution to the dual LP (5.3) for $\text{LP}^{\text{cover}}(\mathcal{S})$. Since $\lambda_{\mathcal{S}}$ are rational numbers, let N be least common multiple of all the $\lambda_{\mathcal{S}}$. We will form a multiset H of minimal open sets in \mathcal{O}^* where each open set $S \in \mathcal{O}^*$ appears in H for $\lambda_S N$ number of times. Let $m = |H| = \sum_S \lambda_S N = N \text{LP}^{\text{cover}}(\mathcal{S})$. Let $\Sigma = \{0, 1\}^m$.

If we write $H = \{F_1, \dots, F_m\}$, then a codeword will be indexed by a tuple of bits $x = (x_1, \dots, x_m) \in \{0, 1\}^m$. To construct the corresponding codeword $w(x) \in \Sigma^n$ we need to specify the value of $w(x)_i \in \Sigma = \{0, 1\}^N$ for each $i = 1, 2, \dots, n$. We set that value to be $w(x)_i = (x_j \mid i \in F_j)$ (ordered in increasing order of j). That is, we associate a bit with each open set in H , and assign the value of a coordinate i to be the list of bits for all sets in H containing i . Note that the alphabet size is bounded by N because $|\{j : i \in F_j\}| = \sum_{S \ni i} \lambda_S N \leq N$.

We now need to show that if $T \models i$ is a rule of \mathcal{S} , then one can recover $w(x)_i$ from $(w(x)_t)_{t \in T}$. Suppose F_j is an open set containing i so that x_j appears in the symbol $w(x)_i$. Then, there must be an element $t \in T$ so that F_j contains t (otherwise, F_j cannot contain i). Thus, the value x_j can be computed from $w(x)_t$. Since this holds for any x_j appearing in $w(x)_i$ we are done.

We will now show that the dimension of a code obtained using Construction 5.4 cannot be better than $\text{LP}^{\text{cover}}(\mathcal{S})$. Let S_1, \dots, S_n be the subsets of size at most ℓ from some universe U obtained from Construction 5.4 such that whenever $T \models i$ in \mathcal{S} , $S_i \subset \cup_{j \in T} S_j$. Such a set system gives a code with dimension $|U|/\ell$. We will show that such a code also gives a feasible solution to the dual LP (see (5.3)). For $u \in U$, let $F_u = \{i : u \in S_i\}$. The complement of F_u is closed since the union of sets which don't have u cannot contain any set which has u ; therefore, F_u is an open set in \mathcal{S} . Let F_u^* be a minimal open set contained in F_u . For a minimal open set $F \in \mathcal{O}^*$, set

$$\lambda_F = \frac{1}{\ell} |\{u \in U : F_u^* = F\}|.$$

Let $i \in [n]$, then

$$\begin{aligned} \sum_{F \in \mathcal{O}^*, F \ni i} \lambda_F &= \sum_{F \in \mathcal{O}^*, F \ni i} \frac{1}{\ell} \sum_{u \in U} \mathbf{1}(F_u^* = F) \\ &= \sum_{u \in U} \sum_{F \in \mathcal{O}^*, F \ni i} \frac{1}{\ell} \mathbf{1}(F_u^* = F) \\ &= \sum_{u \in U} \frac{1}{\ell} \mathbf{1}(i \in F_u^*) \\ &\leq \sum_{u \in U} \frac{1}{\ell} \mathbf{1}(i \in F_u) = \frac{|S_i|}{\ell} = 1. \end{aligned}$$

Thus this assignment is a feasible solution to the dual LP (see (5.3)) with objective value

$$\sum_{F \in \mathcal{O}^*} \lambda_F = \sum_{F \in \mathcal{O}^*} \sum_{u \in U} \frac{1}{\ell} \mathbf{1}(F_u^* = F) = \sum_{u \in U} \sum_{F \in \mathcal{O}^*} \frac{1}{\ell} \mathbf{1}(F_u^* = F) = \sum_{u \in U} \frac{1}{\ell} = \frac{|U|}{\ell}.$$

This implies that $\text{LP}^{\text{cover}}(\mathcal{S}) \geq |U|/\ell$. □

The code construction based on sets achieving dimension equal to $\text{LP}^{\text{cover}}(\mathcal{S})$ constructed in the above lemma needed very large alphabet. The following lemma shows that by using random sampling, one can get very small alphabet and still achieve dimension close to $\text{LP}^{\text{cover}}(\mathcal{S})$.

LEMMA 5.6. *Let \mathcal{S} be a spanoid on $[n]$, then there exists a code $C \subset \Sigma^n$ consistent with \mathcal{S} , obtained as in Construction 5.4 of dimension r over an alphabet Σ such that $\log |\Sigma| \lesssim \log n / \log \log n$ and*

$$r \gtrsim \frac{\log \log n}{\log n} \cdot \text{LP}^{\text{cover}}(\mathcal{S}).$$

Proof. Let us assume that $\text{LP}^{\text{cover}}(\mathcal{S}) \geq \log n / \log \log n$, since otherwise the bound is trivial. Form a random subset $H \subset \mathcal{O}^*$ by sampling each set $S \in \mathcal{O}^*$ with probability λ_S where λ_S is the optimal solution to the dual LP (5.3). Let Z_S be the random variable that S is included in H :

$$\mathbb{E}[|H|] = \mathbb{E} \left[\sum_{S \in \mathcal{O}^*} Z_S \right] = \sum_{S \in \mathcal{O}^*} \lambda_S = \text{LP}^{\text{cover}}(\mathcal{S}),$$

$$\Pr \left[|H| \leq \frac{1}{2} \mathbb{E}[|H|] \right] \leq \exp(-\Omega(\mathbb{E}[|H|])) = o(1).$$

Let $\Delta(H)$ be the maximum number of sets in H that an element of \mathcal{S} belongs to. We want to show that $\Delta(H)$ is small with good probability. Fix some $i \in [n]$. The number of subsets in H which contain i is $\sum_{S \ni i} Z_S$. By Chernoff bound,¹²

$$\Pr \left[\sum_{S \ni i} Z_S \geq t \right] \leq \left(\frac{e \mathbb{E}[\sum_{S \ni i} Z_S]}{t} \right)^t = \left(\frac{e \sum_{S \ni i} \lambda_S}{t} \right)^t \leq \left(\frac{e}{t} \right)^t.$$

By union bound,

$$\Pr[\Delta(H) \geq t] = \Pr \left[\exists i \in [n] : \sum_{S \ni i} Z_S \geq t \right] \leq n \left(\frac{e}{t} \right)^t = o(1)$$

if $t = e \log n / \log \log n$. Therefore, there exists an $H \subset \mathcal{O}^*$ such that $|H| \geq \text{LP}^{\text{cover}}(\mathcal{S})/2$ and $\Delta(H) \leq e \log n / \log \log n$. By imitating the proof of Lemma 5.5 using this H , we get the required result. □

Gaps between LP^{cover} and functional rank. Any code that is obtained by using Construction 5.4 from a q -LCS with error-tolerance δ will have dimension at most $O(q/\delta)$. Indeed, suppose $S_1, \dots, S_n \subset U$ are the sets used in the construction. Then each set S_i is contained in at least $(\delta/q)n$ disjoint q -subsets of sets in $\{S_j : j \neq i\}$ such that $S_i \subset \cup_{j \in T} S_j$ for each q -subset T . Therefore, each element of U occurs in at least (δ/q) fraction of sets in S_1, \dots, S_n . Therefore, a typical set in S_1, \dots, S_n contains at least (δ/q) fraction of elements from U . Therefore, the dimension of such a code can be at most q/δ . This also implies that the LP^{cover} of a q -LCS is $O(q/\delta)$. Since the rank of a q -LCS (with some constant δ) can be $\tilde{\Omega}(n^{1-\frac{1}{q-1}})$ (Theorem 1.10), this

¹²If X_1, X_2, \dots, X_n are Bernoulli random variables with $\mu = \mathbb{E}[\sum_i X_i]$. Then $\Pr[\sum_i X_i \geq t\mu] \leq (e^{t-1}/t^t)^\mu$.

shows that LP^{cover} can be much smaller than the rank of a spanoid. In fact LP^{cover} can be much smaller than the functional rank ($f\text{-rank}$) of a spanoid. $\exp(\tilde{O}(\sqrt{\log n}))$ -query LCCs (for some constant δ) of length n and dimension $\Omega(n)$ are known to exist [KMRS17]. This implies that functional rank of the corresponding spanoid is $\Omega(n)$ whereas LP^{cover} is at most $\exp(\tilde{O}(\sqrt{\log n}))$.

6. Products of spanoids. We have shown that functional rank can be strictly smaller than the rank of a spanoid (Claim 5.3), but it was only a constant factor gap. Can we construct spanoids whose $\text{rank}(\mathcal{S})$ and $f\text{-rank}(\mathcal{S})$ have a polynomial gap? One way to achieve this is by constructing product operations on spanoids which amplify the gap between rank and $f\text{-rank}$. Suppose that given two spanoids \mathcal{S}_1 and \mathcal{S}_2 on $[n_1]$ and $[n_2]$, respectively, we can construct a product spanoid $\mathcal{S}_1 \times \mathcal{S}_2$ on $[n_1] \times [n_2]$, such that $\text{rank}(\mathcal{S}_1 \times \mathcal{S}_2) \geq \text{rank}(\mathcal{S}_1) \cdot \text{rank}(\mathcal{S}_2)$ and $f\text{-rank}(\mathcal{S}_1 \times \mathcal{S}_2) \leq f\text{-rank}(\mathcal{S}_1) \cdot f\text{-rank}(\mathcal{S}_2)$. Then by starting with the pentagon spanoid and taking the above product several times, one can get a spanoid on n elements with $n^{\Omega(1)}$ gap between rank and $f\text{-rank}$. With this motivation, we study a few natural ways to define products of spanoids and study how various notions of rank we defined behave under these products. The following lemma will be useful to compare the rank measures of two spanoids. It shows that a spanoid with more open sets has higher rank , $f\text{-rank}$, LP^{cover} , and $LP^{entropy}$. Intuitively, this is because in the spanoid with more open (closed) sets, you can make fewer inferences.

LEMMA 6.1. *If $\mathcal{S}, \mathcal{S}'$ are two spanoids on X with open sets $\mathcal{O}_{\mathcal{S}}, \mathcal{O}_{\mathcal{S}'}$, respectively. If $\mathcal{O}_{\mathcal{S}} \subset \mathcal{O}_{\mathcal{S}'}$ then*

$$\begin{aligned} \text{rank}(\mathcal{S}) &\leq \text{rank}(\mathcal{S}'), \\ LP^{cover}(\mathcal{S}) &\leq LP^{cover}(\mathcal{S}'), \\ f\text{-rank}(\mathcal{S}) &\leq f\text{-rank}(\mathcal{S}'), \\ LP^{entropy}(\mathcal{S}) &\leq LP^{entropy}(\mathcal{S}'). \end{aligned}$$

Proof. Since $\text{rank}(\mathcal{S})$ is the smallest hitting set for $\mathcal{O}_{\mathcal{S}}$ and $LP^{cover}(\mathcal{S})$ is the smallest fractional hitting set for $\mathcal{O}_{\mathcal{S}}$ and $\mathcal{O}_{\mathcal{S}} \subset \mathcal{O}_{\mathcal{S}'}$, the first two inequalities easily follow. To show the inequality for functional ranks, we will show that if $A \models_{\mathcal{S}'} i$, then $A \models_{\mathcal{S}} i$, i.e., \mathcal{S} has more inference rules than \mathcal{S}' . Thus any code consistent with \mathcal{S} is also consistent with \mathcal{S}' . This also implies the inequality for $LP^{entropy}$ because the corresponding maximization LP for \mathcal{S} has more constraints than for \mathcal{S}' and so the maximum is smaller.

The closed sets $\mathcal{C}_{\mathcal{S}}$ and $\mathcal{C}_{\mathcal{S}'}$ of \mathcal{S} and \mathcal{S}' also satisfy $\mathcal{C}_{\mathcal{S}} \subset \mathcal{C}_{\mathcal{S}'}$. $A \models_{\mathcal{S}'} i$ iff $i \in \text{span}_{\mathcal{S}'}(A)$. By Claim 2.3, $\text{span}_{\mathcal{S}}(A) = \bigcap_{\{B: B \in \mathcal{C}_{\mathcal{S}}, B \supset A\}} B$ and so $\text{span}_{\mathcal{S}}(A) \supset \text{span}_{\mathcal{S}'}(A)$. So if $A \models_{\mathcal{S}'} i$, then $A \models_{\mathcal{S}} i$. \square

6.1. Product $\mathcal{S}_1 \odot \mathcal{S}_2$. In this subsection, we will define a natural product operation where the open sets in the product $\mathcal{S}_1 \odot \mathcal{S}_2$ are products of open sets and their unions.

DEFINITION 6.2. *Given two spanoids $\mathcal{S}_1, \mathcal{S}_2$ on sets X_1, X_2 with collection of open sets $\mathcal{O}_1, \mathcal{O}_2$, respectively, the product spanoid $\mathcal{S}_1 \odot \mathcal{S}_2$ is a spanoid on $X_1 \times X_2$ with open sets given by unions of sets $A \times B$ where $A \in \mathcal{O}_1, B \in \mathcal{O}_2$.*

We will now show that under this product, LP^{cover} is multiplicative.

LEMMA 6.3. $LP^{cover}(\mathcal{S}_1 \odot \mathcal{S}_2) = LP^{cover}(\mathcal{S}_1) \cdot LP^{cover}(\mathcal{S}_2)$.

Proof. Let $\mathcal{S} = \mathcal{S}_1 \odot \mathcal{S}_2$. The minimal open sets in \mathcal{S} are Cartesian products of minimal open sets in \mathcal{S}_1 and \mathcal{S}_2 , i.e.,

$$\mathcal{O}_{\mathcal{S}}^* = \{A \times B : A \in \mathcal{O}_{\mathcal{S}_1}^*, B \in \mathcal{O}_{\mathcal{S}_2}^*\}.$$

$$(6.1) \quad \begin{aligned} \text{LP}^{\text{cover}}(\mathcal{S}) &= \min \sum_{i,j=1}^n z_{ij}, \\ z_{ij} &\geq 0, \\ \sum_{(i,j) \in A \times B} z_{ij} &\geq 1 \quad \forall A \in \mathcal{O}_{\mathcal{S}_1}^*, B \in \mathcal{O}_{\mathcal{S}_2}^*. \end{aligned}$$

Let x, y be the optimal solutions to the LPs (of the form (5.2)) corresponding to $\text{LP}^{\text{cover}}(\mathcal{S}_1)$ and $\text{LP}^{\text{cover}}(\mathcal{S}_2)$, respectively. Then $z = x \otimes y$ (i.e., $z_{ij} = x_i y_j$) is a feasible solution to the above LP for $\text{LP}^{\text{cover}}(\mathcal{S})$. Therefore,

$$\text{LP}^{\text{cover}}(\mathcal{S}) \leq \text{LP}^{\text{cover}}(\mathcal{S}_1) \cdot \text{LP}^{\text{cover}}(\mathcal{S}_2).$$

We can also write the dual LP for $\text{LP}^{\text{cover}}(\mathcal{S})$.

$$(6.2) \quad \begin{aligned} \text{LP}^{\text{cover}}(\mathcal{S}) &= \max \sum_{A \in \mathcal{O}_{\mathcal{S}_1}^*, B \in \mathcal{O}_{\mathcal{S}_2}^*} \lambda_{AB} \\ \lambda_{AB} &\geq 0 \\ \sum_{A \ni i, B \ni j} \lambda_{AB} &\leq 1 \quad \forall i, j \in [n]. \end{aligned}$$

This completes the proof. \square

Let α, β be the optimal solutions to the dual LPs (of the form (5.3)) corresponding to $\text{LP}^{\text{cover}}(\mathcal{S}_1)$ and $\text{LP}^{\text{cover}}(\mathcal{S}_2)$, respectively. Then $\lambda = \alpha \otimes \beta$ (i.e., $\lambda_{AB} = \alpha_A \beta_B$) is a feasible solution to the above dual LP for $\text{LP}^{\text{cover}}(\mathcal{S})$. Therefore,

$$\text{LP}^{\text{cover}}(\mathcal{S}) \geq \text{LP}^{\text{cover}}(\mathcal{S}_1) \cdot \text{LP}^{\text{cover}}(\mathcal{S}_2).$$

6.2. Product $\mathcal{S}_1 \otimes \mathcal{S}_2$. We will now define a different product operation inspired by the following tensor product operation on codes.

DEFINITION 6.4. *Given two codes $C_1 \subset \Sigma^{n_1}$ and $C_2 \subset \Sigma^{n_2}$, the tensor code $C_1 \otimes C_2 \subset \Sigma^{n_1 \times n_2}$ is defined as the set of all $n_1 \times n_2$ matrices over Σ where each column is a codeword in C_1 and each row is a codeword in C_2 .*

If C_1 and C_2 are linear codes over some field \mathbb{F} , i.e., they are linear subspaces of \mathbb{F}^{n_1} and \mathbb{F}^{n_2} , respectively, then the tensor code $C_1 \otimes C_2 \subset \mathbb{F}^{n_1 \times n_2}$ is exactly the tensor product of the subspaces $C_1 \otimes C_2$.

We will now define a product operation on spanoids which mimics the above operation on codes.

DEFINITION 6.5. *Let \mathcal{S}_1 and \mathcal{S}_2 be spanoids on X_1, X_2 , respectively. The product $\mathcal{S}_1 \otimes \mathcal{S}_2$ is a spanoid on $X_1 \times X_2$ generated by the following inference rules:*

1. *For $A \subset X_1, i \in X_1$, if $A \models i$ in \mathcal{S}_1 , then for every $j \in X_2, A \times \{j\} \models (i, j)$.*
2. *For $B \subset X_2, j \in X_2$, if $B \models j$ in \mathcal{S}_2 , then for every $i \in X_1, \{i\} \times B \models (i, j)$.*

The following claim follows easily from the above definitions.

CLAIM 6.6. *Let C_1, C_2 be codes consistent with $\mathcal{S}_1, \mathcal{S}_2$, respectively. Then $C_1 \otimes C_2$ is consistent with $\mathcal{S}_1 \otimes \mathcal{S}_2$.*

How does $\mathcal{S}_1 \otimes \mathcal{S}_2$ compare with $\mathcal{S}_1 \odot \mathcal{S}_2$? The following claim shows that $\mathcal{S}_1 \otimes \mathcal{S}_2$ has more open sets (or closed sets) than $\mathcal{S}_1 \odot \mathcal{S}_2$. This shows that all of the rank measures are smaller for $\mathcal{S}_1 \odot \mathcal{S}_2$ than for $\mathcal{S}_1 \otimes \mathcal{S}_2$.

CLAIM 6.7. $\mathcal{O}_{\mathcal{S}_1 \odot \mathcal{S}_2} \subset \mathcal{O}_{\mathcal{S}_1 \otimes \mathcal{S}_2}$.

Proof. It is enough to show that if $A \in \mathcal{O}_{\mathcal{S}_1}$ and $B \in \mathcal{O}_{\mathcal{S}_2}$, then $A \times B \in \mathcal{O}_{\mathcal{S}_1 \otimes \mathcal{S}_2}$. This is equivalent to showing $(A \times B^c) \cup (A^c \times B)$ being closed in $\mathcal{S}_1 \otimes \mathcal{S}_2$, which is easy to see given the inference rules.

6.3. Product $\mathcal{S}_1 \times \mathcal{S}_2$. In this subsection, we define a product operation called the semidirect product, denoted by $\mathcal{S}_1 \times \mathcal{S}_2$. Under this product we will show that rank is multiplicative and f-rank is submultiplicative. Thus by taking repeated semidirect product of the pentagon spanoid Π_5 with itself, we create a spanoid with polynomial gap between its rank and f-rank which proves Theorem 1.13. Additionally, we also show that $\text{LP}^{\text{entropy}}$ is submultiplicative under this product, which gives a spanoid with a polynomial gap between $\text{LP}^{\text{entropy}}$ and rank. We begin with the formal definition of semidirect product.

DEFINITION 6.8 (semidirect product of spanoids). *Let $\mathcal{S}_1, \mathcal{S}_2$ be two spanoids on X_1, X_2 , respectively. Define $\mathcal{S}_1 \times \mathcal{S}_2$ to be the spanoid on $X_1 \times X_2$ generated by the following rules.*

- (1.) *For $A \subset X_1, i \in X_1$, if $A \models i$ in \mathcal{S}_1 , then for every $j \in X_2, A \times X_2 \models (i, j)$.*
- (2.) *For $B \subset X_2, j \in X_2$, if $B \models j$ in \mathcal{S}_2 , then for every $i \in X_1, \{i\} \times B \models (i, j)$.*

Note how $\mathcal{S}_1 \times \mathcal{S}_2$ differs from $\mathcal{S}_1 \otimes \mathcal{S}_2$ in (1). The semidirect product is not a symmetric product, i.e., $\mathcal{S}_1 \times \mathcal{S}_2$ may not be isomorphic to $\mathcal{S}_2 \times \mathcal{S}_1$. We first show that rank is multiplicative under this product.

LEMMA 6.9. $\text{rank}(\mathcal{S}_1 \times \mathcal{S}_2) = \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$.

Proof. If U_1 is a generating set for \mathcal{S}_1 , and U_2 is a generating set for \mathcal{S}_2 , then $U_1 \times U_2$ is a generating set for $\mathcal{S}_1 \times \mathcal{S}_2$. So $\text{rank}(\mathcal{S}_1 \times \mathcal{S}_2) \leq \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$. We need only prove that $\text{rank}(\mathcal{S}_1 \times \mathcal{S}_2) \geq \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$. The intuition behind this is that, in $\mathcal{S}_1 \times \mathcal{S}_2$, we can w.l.o.g. assume that all the \mathcal{S}_2 -derivations (i.e., rule (2) in Definition 6.8) are done before all the \mathcal{S}_1 -derivations (i.e., rule (1) in Definition 6.8). We will write this more formally below.

Let T be a generating set for $\mathcal{S}_1 \times \mathcal{S}_2$. We would like to prove that $|T| \geq \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$. Let $T_0 = T, T_1, \dots, T_r = T_1 \times T_2$, and A_1, \dots, A_r be such that for all $k \in [r]$, we have $T_{k-1} \subsetneq T_k, A_k \subset T_{k-1}$, and one of the following is true:

1. $T_k = T_{k-1} \cup (\{i\} \times X_2)$ for some $i \in X_1, A_k = A \times X_2$ for some $A \subset X_1$, and $A \models i$ in \mathcal{S}_1 .
2. We are not in case (1), and $T_k = T_{k-1} \cup \{(i, j)\}$ for some $(i, j) \in X_1 \times X_2, A_k = \{i\} \times B$ for some $B \subset X_2$, and $B \models j$ in \mathcal{S}_2 .

If case (1) is true, we say step k is a \mathcal{S}_1 -derivation. If case (2) is true, we say step k is a \mathcal{S}_2 -derivation.

CLAIM 6.10. *We can choose $T_0, \dots, T_r, A_1, \dots, A_r$ such that there exists an integer l for which*

1. *for all $k \leq l$, step k is a \mathcal{S}_2 -derivation;*
2. *for all $k \geq l + 1$, step k is a \mathcal{S}_1 -derivation.*

Proof. Suppose there exists some k such that step k is a \mathcal{S}_1 -derivation and step $k + 1$ is a \mathcal{S}_2 -derivation. It is not hard to see that we can swap step k and step $k + 1$. Repeatedly applying this until no such k exists, and we get the desired sequences. \square

Now we return to the proof of $|T| \geq \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$. Let l be the integer in the above claim. Because steps $k \geq l + 1$ are all of type (1), we have

$$|\{i : \{i\} \times X_2 \subset T_l\}| \geq \text{rank}(\mathcal{S}_1).$$

Because steps $k \leq l$ are all of type (2), for each i such that $\{i\} \times X_2 \subset T_l$, we have

$$|T \cap (\{i\} \times X_2)| \geq \text{rank}(\mathcal{S}_2).$$

Combining the two inequalities we get $|T| \geq \text{rank}(\mathcal{S}_1)\text{rank}(\mathcal{S}_2)$. \square

We now show that f -rank is submultiplicative under semidirect product.

LEMMA 6.11. $f\text{-rank}(\mathcal{S}_1 \times \mathcal{S}_2) \leq f\text{-rank}(\mathcal{S}_1)f\text{-rank}(\mathcal{S}_2)$.

Proof. Let $\mathcal{C} \in \Sigma^{X_1 \times X_2}$ be a code consistent with $\mathcal{S}_1 \times \mathcal{S}_2$. For $i \in X_1$, define code $\mathcal{C}_2^i \subset \Sigma^{X_2}$ as $\mathcal{C}_2^i = \{c_{\{i\} \times X_2} : c \in \mathcal{C}\}$. Because \mathcal{C}_2^i is consistent with \mathcal{S}_2 , we have $\log |\mathcal{C}_2^i| / \log |\Sigma| \leq f\text{-rank}(\mathcal{S}_2)$. Let $N = \max_{i \in X_1} |\mathcal{C}_2^i|$ and for each $i \in X_1$, choose an injection $\phi_i : \mathcal{C}_2^i \hookrightarrow [N]$. Clearly, we have

$$\frac{\log N}{\log |\Sigma|} \leq f\text{-rank}(\mathcal{S}_2).$$

Define $\mathcal{C}_1 \subset [N]^{X_1}$ as $\mathcal{C}_1 = \{(\phi_i(c_{\{i\} \times X_2}))_{i \in X_1} : c \in \mathcal{C}\}$. Then $|\mathcal{C}_1| = |\mathcal{C}|$ and \mathcal{C}_1 is a code consistent with \mathcal{S}_1 . So we have

$$\frac{\log |\mathcal{C}_1|}{\log N} \leq f\text{-rank}(\mathcal{S}_1).$$

Combining the inequalities, we get

$$\frac{\log |\mathcal{C}|}{\log |\Sigma|} \leq f\text{-rank}(\mathcal{S}_1)f\text{-rank}(\mathcal{S}_2).$$

We can now prove Theorem 1.13.

Proof of Theorem 1.13. Define $\mathcal{S}_1 = \Pi_5$, and $\mathcal{S}_i = \Pi_5 \times \mathcal{S}_{i-1}$ for $i \geq 2$. Then \mathcal{S}_n is a spanoid on 5^n elements. By Lemma 6.11,

$$f\text{-rank}(\mathcal{S}_n) \leq f\text{-rank}(\Pi_5)^n = 2.5^n.$$

By Lemma 6.9,

$$\text{rank}(\mathcal{S}_n) = \text{rank}(\Pi_5)^n = 3^n.$$

So

$$\text{rank}(\mathcal{S}_n) \geq (5^n)^{\log_5 3 - \log_5 2.5} f\text{-rank}(\mathcal{S}_n).$$

We also show that $\text{LP}^{\text{entropy}}$ is submultiplicative under semidirect product.

LEMMA 6.12. $\text{LP}^{\text{entropy}}(\mathcal{S}_1 \times \mathcal{S}_2) \leq \text{LP}^{\text{entropy}}(\mathcal{S}_1)\text{LP}^{\text{entropy}}(\mathcal{S}_2)$.

Proof. Recall the linear program in (5.1) that defines $\text{LP}^{\text{entropy}}$. Let $f : 2^{X_1 \times X_2} \rightarrow \mathbb{R}_{\geq 0}$ be the optimal solution of the linear program which computes $\text{LP}^{\text{entropy}}(\mathcal{S}_1 \times \mathcal{S}_2)$, and so $f(X_1 \times X_2) = \text{LP}^{\text{entropy}}(\mathcal{S}_1 \times \mathcal{S}_2)$. Define $f_1 : 2^{X_1} \rightarrow \mathbb{R}_{\geq 0}$ as

$$f_1(A) = \frac{f(A \times X_2)}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)}.$$

We claim that f_1 satisfies the linear program for \mathcal{S}_1 . We will check all the feasibility conditions of the LP given in (5.1):

1. $f_1(\emptyset) = \frac{f(\emptyset)}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} = 0$.
2. For all $i \in X_1$, the function $f_2^i : 2^{X_2} \rightarrow \mathbb{R}_{\geq 0}$, defined as $f_2^i(B) = f(\{i\} \times B)$, is a feasible solution for the LP which computes $\text{LP}^{\text{entropy}}(\mathcal{S}_2)$, which is a maximization LP. So

$$f_1(\{i\}) = \frac{f_2^i(X_2)}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} \leq 1.$$

3. For all $A, B \subset X_1$, we have

$$\begin{aligned} & f_1(A \cup B) + f_1(A \cap B) \\ &= \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} (f((A \cup B) \times X_2) + f((A \cap B) \times X_2)) \\ &= \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} (f((A \times X_2) \cup (B \times X_2)) + f((A \times X_2) \cap (B \times X_2))) \\ & \text{(submodularity of } f) \\ &\leq \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} (f(A \times X_2) + f(B \times X_2)) \\ &= f_1(A) + f_1(B). \end{aligned}$$

4. For $A \subset B \subset X_1$, we have

$$f_1(A) = \frac{f(A \times X_2)}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} \leq \frac{f(B \times X_2)}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} = f_1(B).$$

5. Let $A \models i$ in \mathcal{S}_1 . Then $A \times X_2 \models (i, j)$ for all $j \in X_2$. So

$$\begin{aligned} f_1(A \cup \{i\}) &= \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} f((A \cup \{i\}) \times X_2) \\ &= \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} f((A \times X_2) \cup (\{i\} \times X_2)) \\ &= \frac{1}{\text{LP}^{\text{entropy}}(\mathcal{S}_2)} f(A \times X_2) \\ &= f_1(A). \end{aligned} \quad \square$$

So f_1 is a feasible solution for the linear program which computes $\text{LP}^{\text{entropy}}(\mathcal{S}_1)$, which is a maximization LP. Therefore, $f_1(X_1) \leq \text{LP}^{\text{entropy}}(\mathcal{S}_1)$. So

$$\text{LP}^{\text{entropy}}(\mathcal{S}_1 \ltimes \mathcal{S}_2) = f(X_1 \times X_2) = f_1(X_1) \text{LP}^{\text{entropy}}(\mathcal{S}_2) \leq \text{LP}^{\text{entropy}}(\mathcal{S}_1) \text{LP}^{\text{entropy}}(\mathcal{S}_2).$$

Thus by repeatedly taking semidirect product of the pentagon spanoid Π_5 with itself, we can obtain a polynomial gap between $\text{LP}^{\text{entropy}}$ and rank , which is a stronger statement than Theorem 1.13 because $\text{f-rank}(\mathcal{S}) \leq \text{LP}^{\text{entropy}}(\mathcal{S})$.

COROLLARY 6.13. *There exists a spanoid \mathcal{S} on n elements with $\text{rank}(\mathcal{S}) \geq n^c \text{LP}^{\text{entropy}}(\mathcal{S})$, where $c = \log_5 3 - \log_5 2.5 \geq 0.113$.*

Because $\mathcal{S}_1 \ltimes \mathcal{S}_2$ has fewer inferences than $\mathcal{S}_1 \otimes \mathcal{S}_2$, every closed set in $\mathcal{S}_1 \otimes \mathcal{S}_2$ is also closed in $\mathcal{S}_1 \ltimes \mathcal{S}_2$ i.e., $\mathcal{O}_{\mathcal{S}_1 \otimes \mathcal{S}_2} \subset \mathcal{O}_{\mathcal{S}_1 \ltimes \mathcal{S}_2}$. Therefore, we have the following relationship between the different products we constructed:

$$(6.3) \quad \mathcal{O}_{\mathcal{S}_1 \odot \mathcal{S}_2} \subset \mathcal{O}_{\mathcal{S}_1 \otimes \mathcal{S}_2} \subset \mathcal{O}_{\mathcal{S}_1 \ltimes \mathcal{S}_2}.$$

By Lemma 6.1, all of the rank measures should follow the same order. Therefore, we have the following corollary.

COROLLARY 6.14. *For any two spanoids \mathcal{S}_1 and \mathcal{S}_2 ,*

$$\begin{aligned} \text{rank}(\mathcal{S}_1 \odot \mathcal{S}_2) &\leq \text{rank}(\mathcal{S}_1 \otimes \mathcal{S}_2) \leq \text{rank}(\mathcal{S}_1 \times \mathcal{S}_2) = \text{rank}(\mathcal{S}_1) \cdot \text{rank}(\mathcal{S}_2) \\ LP^{\text{entropy}}(\mathcal{S}_1 \odot \mathcal{S}_2) &\leq LP^{\text{entropy}}(\mathcal{S}_1 \otimes \mathcal{S}_2) \leq LP^{\text{entropy}}(\mathcal{S}_1 \times \mathcal{S}_2) \\ &\leq LP^{\text{entropy}}(\mathcal{S}_1) \cdot LP^{\text{entropy}}(\mathcal{S}_2) \\ f\text{-rank}(\mathcal{S}_1 \odot \mathcal{S}_2) &\leq f\text{-rank}(\mathcal{S}_1 \otimes \mathcal{S}_2) \leq f\text{-rank}(\mathcal{S}_1 \times \mathcal{S}_2) \leq f\text{-rank}(\mathcal{S}_1) \cdot f\text{-rank}(\mathcal{S}_2) \\ LP^{\text{cover}}(\mathcal{S}_1) \cdot LP^{\text{cover}}(\mathcal{S}_2) &= LP^{\text{cover}}(\mathcal{S}_1 \odot \mathcal{S}_2) \leq LP^{\text{cover}}(\mathcal{S}_1 \otimes \mathcal{S}_2) \leq LP^{\text{cover}}(\mathcal{S}_1 \times \mathcal{S}_2). \end{aligned}$$

Thus LP^{cover} is supermultiplicative under semidirect product, while rank is multiplicative and the other two are submultiplicative.

Example 6.15. In contrast to the semidirect product, rank is not multiplicative under $\mathcal{S}_1 \odot \mathcal{S}_2$ and $\mathcal{S}_1 \otimes \mathcal{S}_2$.

- This example is found by Xu [Xu18]. Let \mathcal{S} be the spanoid on a four element set $\{1, 2, 3, 4\}$ with the following closed sets (using the equivalence with intersection closed families):

$$\emptyset, \{1\}, \{4\}, \{2, 4\}, \{3, 4\}, \{1, 2, 3, 4\}.$$

Clearly, $\text{rank}(\mathcal{S}) = 2$. On the other hand, the set $\{(1, 1), (2, 3), (3, 2)\}$ generates $\mathcal{S}_1 \otimes \mathcal{S}_2$. So $\text{rank}(\mathcal{S}_1 \otimes \mathcal{S}_2) \leq 3 < 4 = \text{rank}(\mathcal{S})^2$.

- In fact, the spanoid $\Pi_5 \otimes \Pi_5$ admits a generating set of size 8, and therefore, $\text{rank}(\Pi_5 \otimes \Pi_5) \leq 8 < 9 = \text{rank}(\Pi_5)^2$. One such generating set is

$$\{(1, 1), (1, 2), (2, 1), (2, 2), (3, 3), (3, 4), (4, 3), (5, 5)\}.$$

7. Conclusion and open problems. Our work introduces the abstract notion of a spanoid in the hope that further study of its properties will lead to progress on LCCs and perhaps in other areas. We list below some concrete directions for future work.

1. We showed that there exist spanoids, called q -LCSs, which “look like” q -LCCs and whose rank matches the best known upper bounds. Can we bypass this “barrier” by using additional properties of LCCs? We have at least two examples where this was possible. One is the result of [KW04] for LCCs over constant size alphabet and the other is the work in [DSW14] for linear 3-LCCs over the real numbers. The bounds of [KW04] crucially depend on the alphabet having small size and the bounds in [DSW14] exploit properties of real numbers.
2. Understanding the possible gap between functional rank and formal rank of a spanoid is a very interesting question. We proved that there can be a polynomial gap. The next challenge is to find a spanoid on n elements whose $f\text{-rank}$ is $n^{o(1)}$ and rank is $n^{\Omega(1)}$. Naturally, q -LCSs for constant $q \geq 3$ are plausible candidates for this. If there are no such spanoids, then it would imply the existence of q -LCCs of length n and $n^{\Omega_q(1)}$ dimension!¹³
3. Are there general methods (in the spirit of Construction 5.4, which we show is limited) to achieve high functional rank?

¹³Possibly over a large alphabet.

4. We have constructed spanoid products under which rank is multiplicative or LP^{cover} is multiplicative. Can we construct a spanoid product under which f-rank is multiplicative?
5. Suppose we start with a functional representation with large alphabet, can we do alphabet reduction without losing too many codewords?
6. We have seen that one way to go past the rank barrier is to use $\text{LP}^{\text{entropy}}$. Can we improve the existing upper bounds on the dimension of q -LCCs by upper bounding $\text{LP}^{\text{entropy}}$ of q -LCSs? Can we use LP duality and construct good feasible solutions to the dual of $\text{LP}^{\text{entropy}}$ to prove good upper bounds on $\text{LP}^{\text{entropy}}$?
7. For a spanoid \mathcal{S} arising from a matroid, $\text{LP}^{\text{entropy}}(\mathcal{S}) = \text{rank}(\mathcal{S})$. This is because the rank function of a matroid is a feasible solution to the LP (5.1). Can we separate rank and f-rank for spanoids arising from matroids? One possibility is to use non-Shannon type information inequalities.
8. What are other connections of spanoids to existing theory of set systems, matroids, algebraic equations, and other problems described in the introduction?

Acknowledgments. The second author would like to thank Sumegha Garg for helpful discussions. The third author would like to thank Yury Polyanskiy for helpful discussions.

REFERENCES

- [ABRW04] M. ALEKHNIVICH, E. BEN-SASSON, A. A. RAZBOROV, AND A. WIGDERSON, *Pseudorandom generators in propositional proof complexity*, SIAM J. Comput., 34 (2004), pp. 67–88.
- [AR01] M. ALEKHNIVICH AND A. A. RAZBOROV, *Lower bounds for polynomial calculus: Non-binomial case*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, 2001, IEEE, 2001, pp. 190–199.
- [BDSS11] A. BHATTACHARYYA, Z. DVIR, A. SHPILKA, AND S. SARAF, *Tight lower bounds for 2-query lcs over finite fields*, in Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science. IEEE, 2011, pp. 638–647.
- [BDWY12] B. BARAK, Z. DVIR, A. WIGDERSON, AND A. YEHUDAYOFF, *Fractional Sylvester-Gallai theorems*, Proceedings of the National Academy of Sciences, 2012.
- [BDYW11] B. BARAK, Z. DVIR, A. YEHUDAYOFF, AND A. WIGDERSON, *Rank bounds for design matrices with applications to combinatorial geometry and locally correctable codes*, in Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, ACM, 2011, pp. 519–528.
- [BG95] H. BRÖNNIMANN AND M. T. GOODRICH, *Almost optimal set covers in finite vc -dimension*, Discrete Computat. Geom., 14 (1995), pp. 463–479.
- [BGT17] A. BHATTACHARYYA, S. GOPI, AND A. TAL, *Lower bounds for 2-query LCCs over large alphabet*, in Proceedings of the Approximation, Randomization, and Combinatorial Optimization, Algorithms and Techniques, APPROX/RANDOM 2017, 2017, Berkeley, CA, 2017, pp. 30:1–30:20.
- [Bol68] B. BOLLOBÁS, *Weakly k -saturated graphs*, Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967), Teubner, Leipzig, 1968, pp. 25–31.
- [BS15] H. BRUHN AND O. SCHAUDT, *The journey of the union-closed sets conjecture*, Graphs Combin., 31 (2015), pp. 2043–2074 .
- [BW01] E. BEN-SASSON AND A. WIGDERSON, *Short proofs are narrow-resolution made simple*, J. ACM, 48 (2001), pp. 149–169.
- [CT91] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, Wiley Series in Telecommunications. John Wiley and Sons, New York, 1991.
- [DK11] I. DINUR AND T. KAUFMAN, *Dense locally testable codes cannot have constant rate and distance*, in Proceedings of the Approximation, Randomization, and Combinatorial Optimization, Algorithms and Techniques Springer, 2011, pp. 507–518.
- [DSW14] Z. DVIR, S. SARAF, AND A. WIGDERSON, *Breaking the quadratic barrier for 3-LCC's*

- over the reals, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing, ACM, 2014, pp. 784–793.
- [Dvi12] Z. DVIR, *Incidence theorems and their applications*, Found. Trends Theor. Comput. Sci., 6(2012), pp. 257–393.
- [EFF85] P. ERDŐS, P. FRANKL, AND Z. FÜREDI, *Families of finite sets in which no set is covered by the union of others*, Israel J. Math., 51 (1985), pp. 79–89.
- [Efr09] K. EFREMENKO, *3-query locally decodable codes of subexponential length*, in Proceedings of the 2009 ACM International Symposium on Theory of Computing, ACM; New York, 2009, pp. 39–44.
- [ERS05] G. EVEN, D. RAWITZ, AND S. M. SHAHAR, *Hitting sets when the vc-dimension is small*, Inform. Process. Lett., 95 (2005), pp. 358–362.
- [F96] Z. FÜREDI, *ONR-cover-free families*, J. Combin. Theory Ser. A, 73 (1996), pp. 172–173.
- [FLS12] J. FOX, C. LEE, AND B. SUDAKOV, *Maximum union-free subfamilies*, Israel J. Math., 191 (2012), pp. 959–971.
- [HHQ17] L. HAMBARDZUMYAN, H. HATAMI, AND Y. QIAN, *Polynomial Method and Graph Bootstrap Percolation*, preprint, <https://arxiv.org/abs/1708.04640>, 2017.
- [HLW06] S. HOORY, N. LINIAL, AND A. WIGDERSON, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.), 43 (2006), pp. 439–561.
- [IS18] E. ICELAND AND A. SAMORODNITSKY, *On Coset Leader Graphs of Structured Linear Codes*, preprint, <https://arxiv.org/abs/1802.01184>, 2018.
- [Kel86] L. M. KELLY, *A resolution of the Sylvester-Gallai problem of J.-P. Serre*, Discrete Comput. Geom., 1 (1986), pp. 101–104.
- [KMRS17] S. KOPPARTY, O. MEIR, N. RON-ZEWI, AND S. SARAF, *High-rate locally correctable and locally testable codes with sub-polynomial query complexity*, J. ACM, 64 (2017), 11.
- [Kni94] E. KNILL, *Graph Generated Union-closed Families of Sets*, preprint, <https://arxiv.org/abs/math/9409215>, 1994.
- [KT00] J. KATZ AND L. TREVISAN, *On the efficiency of local decoding procedures for error-correcting codes*, in Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC 2000), ACM, 2000, pp. 80–86.
- [KW04] I. KERENIDIS AND R. DE WOLF, *Exponential lower bound for 2-query locally decodable codes via a quantum argument*, J. Comput. System Sci., 69 (2004), pp. 395–420, preliminary version appeared in STOC’03.
- [Mat07] F. MATUS, *Infinitely many information inequalities*, in Proceedings of the IEEE International Symposium on Information Theory, 2007, ISIT 2007, IEEE, 2007, pp. 41–44.
- [Shp18] A. SHPILKA, *Sylvester-Gallai Type Theorems for Quadratic Polynomials*, private communication, 2018.
- [Wój99] P. WÓJCIK, *Union-closed families of sets*, Discrete Math., 199 (1999), pp. 173–182.
- [Woo07] D. WOODRUFF, *New lower bounds for general locally decodable codes*, in Electronic Colloquium on Computational Complexity (ECCC), 14, 2007, TR07-006.
- [Xu18] Y. XU, private communication, 2018.
- [Yek08] S. YEKHANIN, *Towards 3-query locally decodable codes of subexponential length*, J. ACM, 55 (2008), 1.
- [Yek12] S. YEKHANIN, *Locally decodable codes*, Found. Trends Theor. Comput. Sci., 6 (2012), pp. 139–255.
- [Yeu08] R. W. YEUNG, *Information Theory and Network Coding*, Springer-Verlag, New York, 2008.
- [ZY98] Z. ZHANG AND R. W. YEUNG, *On characterization of entropy function via information inequalities*, IEEE Trans. Inform. Theory, 44 (1998), pp. 1440–1452.