Lexicographic Products and the Power of Non-Linear Network Coding

Anna Blasiak* Robert Kleinberg† Eyal Lubetzky‡

Abstract

We introduce a technique for establishing and amplifying gaps between parameters of network coding and index coding problems. The technique uses linear programs to establish separations between combinatorial and coding-theoretic parameters and applies hypergraph lexicographic products to amplify these separations. This entails combining the dual solutions of the lexicographic multiplicands and proving that this is a valid dual solution of the product. Our result is general enough to apply to a large family of linear programs. This blend of linear programs and lexicographic products gives a recipe for constructing hard instances in which the gap between combinatorial or coding-theoretic parameters is *polynomially large*. We find polynomial gaps in cases in which the largest previously known gaps were only small constant factors or entirely unknown. Most notably, we show a polynomial separation between linear and non-linear network coding rates. This involves exploiting a connection between matroids and index coding to establish a previously unknown separation between linear and non-linear index coding rates. We also construct index coding problems with a polynomial gap between the broadcast rate and the trivial lower bound for which no gap was previously known.

1 Introduction

The problem of *Network Coding*, introduced by Ahlswede *et al* [2] in 2000, asks for the maximum rate at which information can be passed from a set of sources to a set of targets in a capacitated network. In practice, there are many examples where network coding provides faster transmission rates compared to traditional routing, e.g. [8] details a recent one in wireless networks. However, despite tremendous initial success in using network coding to solve some *broadcast* problems (those in which every receiver demands the same message), very little is known about how to compute or approximate the network coding rate in general. (See [12] for a survey of the topic.)

In the absence of general algorithms for solving network coding, attention has naturally turned to restricted models of coding (e.g. linear functions between vector spaces over finite fields) and to approximating network coding rates using graph-theoretic parameters (e.g. minimum cut [1] and the independence number [3]). Several of these variants provide bounds on the network coding

*Department of Computer Science, Cornell University, Ithaca NY 14853. E-mail: ablasiak@cs.cornell.edu. Supported by an NDSEG Graduate Fellowship, an AT&T Labs Graduate Fellowship, and an NSF Graduate Fellowship.

†Department of Computer Science, Cornell University, Ithaca NY 14853. E-mail: rdk@cs.cornell.edu. Supported in part by NSF grant CCF-0729102, AFOSR grant FA9550-09-1-0100, a Microsoft Research New Faculty Fellowship, a Google Research Grant, and an Alfred P. Sloan Foundation Fellowship.

‡Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. Email: eyal@microsoft.com.

rate, but the worst-case approximation factor of these bounds remains unknown. For example, it is known that there exists a network in which non-linear network coding can achieve a rate which exceeds the best linear network code by a factor of $\frac{11}{10}$ [6], but it is not known whether this gap[1] can be improved to $n^{1-\varepsilon}$, or even possibly to $\Theta(n)$.

In this paper we introduce a general technique for amplifying many of these gaps by combining linear programming with hypergraph product operations. For instance, this enables us to construct a family of network coding instances with n messages, in which the rate of the best non-linear network code exceeds the rate of the best (vector-)linear network code by a factor of at least n^{ε}. A crucial ingredient in our technique is *index coding* [4, 5], a class of communication problems in which a server holds a set of messages that it wishes to broadcast over a noiseless channel to a set of receivers. Each receiver is interested in one of the messages and has side-information comprising some subset of the other messages. The objective is to devise an optimal encoding scheme (one minimizing the broadcast length) that allows all the receivers to retrieve their required information. Following [3], we use β to denote the limiting value of the information rate (i.e., ratio of broadcast length to message length) of this optimal scheme, as the message length tends to infinity.

In our framework, index coding is most useful for isolating a sub-class of network coding problems that can be combined using lexicographic products. However, it is also an important and well-studied problem in its own right. Index coding is intimately related to network coding in general. It is essentially equivalent to the special case of network coding in which only one edge has finite capacity.[2] Additionally, [11] shows that linear network coding can be reduced to linear index coding, thus implying that index coding captures much of the difficulty of network coding.

Index coding is also intricately related to other well-studied areas of mathematics. Connections between matroids and index coding were established in [10]; for example, that paper shows that realizability of a matroid over a field $\mathbb{F}$ is equivalent to linear solvability of a corresponding index coding problem. Index coding is also closely connected to graph theory: a special case of index coding can be described by an undirected graph G, representing a communication problem where a broadcast channel communicates messages to a set of vertices, each of whom has side-information consisting of the neighbors' messages. Letting $\alpha(G), \overline{\chi}(G)$ denote the independence and clique-cover numbers of G, respectively, one has

$$\alpha(G) \leq \beta(G) \leq \overline{\chi}(G)\,. \tag{1.1}$$

The first inequality above is due to an independent set being identified with a set of receivers with no mutual information, whereas the last one due to [4, 5] is obtained by broadcasting the bitwise XOR of the vertices per clique in the optimal clique-cover of G. As one consequence of the general technique we develop here, we settle an open question of [3] by proving that $\alpha(G)$ can differ from $\beta(G)$; indeed, we show that their ratio can be as large as $n^{0.139}$.

[1] The literature on network coding distinguishes between *linear* network codes, in which the messages are required to be elements of a finite field, and *vector-linear* network codes, in which the messages are elements of a finite-dimensional vector space over a finite field. Linear coding is weaker, and a gap of size $n^{1-\varepsilon}$ is known [9]. Vector-linear coding is much more powerful, and no gap larger than 11/10 was known prior to our work.

[2] The unique finite-capacity edge represents the broadcast channel. Each sender is connected to the tail of this edge, each receiver is connected to its head, and each receiver has incoming edges directly from a subset of the senders, representing the side-information.

1.1 Contributions

We present a general technique that amplifies lower bounds for index coding problems using lexicographic hypergraph products in conjunction with linear programs that express information-theoretic inequalities. The use of such linear programs to prove lower bounds in network coding theory is not new, but, perhaps surprisingly, they have not gained widespread use in the analysis of index coding problems. We give an information-theoretic linear program, whose solution, b, gives the best known lower bound on β. However, our main innovation is the insight that this linear programming technique can be combined with the combinatorial technique of graph products to yield lower bounds for sequences of index coding and network coding problems. Specifically, we provide a lexicographic product operation on index coding problems along with an operation that combines dual solutions of the corresponding two linear programs. We show that the combined dual yields a dual solution of the linear program corresponding to the lexicographic product. Using this operation, we demonstrate that index coding lower bounds proven using linear programming behave supermultiplicatively under lexicographic products. This technical tool enables us to prove some new separation results answering open questions in the field.

Our technique not only applies to the standard linear programs used in network information theory (those that express entropy inequalities such as submodularity) but to *any* family of linear programs constructed using what we call a *tight homomorphic constraint schema*. In particular, if one can develop a tight homomorphic constraint schema that applies to a restricted class of codes (e.g. linear) then it becomes possible to prove lower bounds for this class of codes and amplify them using lexicographic products. We pursue this approach in establishing a large multiplicative gap between linear and non-linear network coding.

Theorem 1.1. *Lower bounds for index coding problems can be proven by solving a linear program whose constraints are valid for the class of coding functions being considered. If the linear program is constructed using a tight homomorphic constraint schema (see Section 3), then its optimum is supermultiplicative under the lexicographic product of two index coding problems.*

To separate linear from non-linear coding, we first produce a pair of linear inequalities that are valid information inequalities for tuples of random variables defined by linear functions over fields of odd (resp., even) characteristic, but not vice-versa. We obtain these inequalities by considering the Fano and non-Fano matroids; the former is a matroid that is only realizable in characteristic 2, while the latter is only realizable in odd characteristic and in characteristic 0. For each of the two matroids, we are able to transform a proof of its non-realizability into a much stronger quantitative statement about dimensions of vector spaces over a finite field. This, in turn, we transform into a tight homomorphic constraint schema of valid information inequalities for linear random variables.

We then use the connection between matroids and index coding [6, 7, 10] and these inequalities to give a pair of index coding instances where the best non-linear coding rate is strictly better than the best linear rate over a field of odd (resp.,even) characteristic. We do this by establishing a general theorem that says that for a matroid M, and an inequality that is violated for the rank function of M, there is an index coding problem for which the bound obtained by adding this inequality to the LP is strictly greater than b.

We can now plug the constraint schema into our lexicographic product technique and apply it to these two index coding problems to yield the aforementioned separation between (vector-)linear and non-linear network coding.

Theorem 1.2. *There exists an explicit family of network coding instances (based on index coding instances) with n messages and some fixed $\varepsilon > 0$ such that the non-linear rate is $\Omega(n^{\varepsilon})$ times larger than the linear rate.*

The largest previously known gap between the non-linear and linear rates for network coding was a factor of $\frac{11}{10}$ ([7]). No separation was known between these parameters for index coding (see [3, 9] for related separation results focusing on the weaker setting of scalar linear codes).

As explained above, given any index coding problem G we can write down an LP whose constraints are based on information inequalities that gives a lower bound on β. It is the best known lower bound, and in many cases, strictly better than any previously known bound. Notably, we can show that the broadcast rate of the 5-cycle is at least $\frac{5}{2}$, giving the first known gap between the independence number α (which equals 2 for the 5-cycle) and the broadcast rate β. Amplifying this gap using lexicographic products, we can boost the ratio β/α to grow polynomially with n in a family of n-vertex graphs.

Theorem 1.3. *There exists an explicit family of index coding instances with n messages such that $\beta(G)$ is at least $\Omega(n^{\delta})$ times larger than $\alpha(G)$, where $\delta = 1 - 2\log_5(2) \approx 0.139$.*

The remainder of the paper is organized as follows. In Section 2 we give a formal definition of index coding and the lexicographic product of two index coding problems. In Section 3 we describe a general class of LPs and prove they behave supermultiplicatively under lexicographic products. Section 4 is devoted to the proof of Theorem 1.3. In Section 5 we give a construction from matroids to index coding and prove a number of connections between properties of the matroid and the parameters of the corresponding index coding problem. Finally, in Section 6 we establish inequalities that are valid for linear codes over fields of odd (resp., even) characteristic and then use these to prove Theorem 1.2.

2 Definitions

An index coding problem is specified by a *directed hypergraph* $G = (V, E)$, where elements of V are thought of as messages, and $E \subseteq V \times 2^V$ is a set of directed hyperedges (v, S), each of which is interpreted as a receiver who already knows the messages in set S and wants to receive message v. Messages are drawn from a finite alphabet Σ, and a solution of the problem specifies a finite alphabet Σ_P to be used by the public channel, together with an encoding scheme $\mathcal{E} : \Sigma^{|V|} \to \Sigma_P$ such that, for any possible values of $(x_v)_{v \in V}$, every receiver (v, S) is able to decode the message x_v from the value of $\mathcal{E}(\vec{x})$ together with that receiver's side information. The minimum encoding length $\ell = \lceil \log_2 |\Sigma_P| \rceil$ for messages that are t bits long (i.e. $\Sigma = \{0,1\}^t$) is denoted by $\beta_t(G)$. As noted in [9], due to the overhead associated with relaying the side-information map to the server the main focus is on the case $t \gg 1$ and namely on the following *broadcast rate*.

$$\beta(G) \triangleq \lim_{t \to \infty} \frac{\beta_t(G)}{t} = \inf_t \frac{\beta_t(G)}{t} \tag{2.1}$$

(The limit exists by subadditivity.) This is interpreted as the average asymptotic number of broadcast bits needed per bit of input, that is, the asymptotic broadcast rate for long messages. We are also interested in the optimal rate when we require that Σ is a finite-dimensional vector space over a finite field $\mathbb{F}$, and the encoding function is linear. We denote this by $\lambda^{\mathbb{F}}$, and we denote the optimal linear rate over any field as λ.

A useful notion in index coding is the following *closure* operation with respect to G, a given instance of the problem: for a set of messages $S \subseteq V$, define

$$\mathrm{cl}(S) = \mathrm{cl}_G(S) = S \cup \{x \mid \exists (x, T) \in E \text{ s.t. } T \subseteq S\}. \tag{2.2}$$

The interpretation is that every message $x \in \mathrm{cl}(S)$ can be decoded by someone who knows all of the messages in S in addition to the broadcast message. In Section 5 when we discuss a transformation that associates an index coding problem to every matroid, the closure operation defined in this paragraph — when specialized to the index coding problems resulting from that transformation — will coincide with the usual matroid-theoretic closure operation.

We next define the lexicographic product operation for directed hypergraphs, then proceed to present Theorem 2.2 which demonstrates its merit in the context of index coding by showing that β is submultiplicative for this operation. The proof gives further intuition for the product operation.

Definition 2.1. The lexicographic product of two directed hypergraphs G, F, denoted by $G \bullet F$, is a directed hypergraph whose vertex set is the cartesian product $V(G) \times V(F)$. The edge set of $G \bullet F$ contains a directed hyperedge e for every pair of hyperedges $(e_G, e_F) \in E(G) \times E(F)$. If $e_G = (w_G, S_G)$ and $e_F = (w_F, S_F)$, then the head of $e = (e_G, e_F)$ is the ordered pair (w_G, w_F) and its tail is the set $(S_G \times V(F)) \cup (\{w_G\} \times S_F)$. Denote by $G^{\bullet n}$ the n-fold lexicographic power of G.

Remark. In the special case where the index coding problem is defined by a graph[3] the above definition coincides with the usual lexicographic graph product (where $G \bullet F$ has the vertex set $V(G) \times V(F)$ and an edge from (u, v) to (u', v') iff either $(u, u') \in E(G)$ or $u = u'$ and $(v, v') \in E(F)$).

Theorem 2.2. *The broadcast rate is submultiplicative under the lexicographic product of index coding problems. That is, $\beta(G \bullet F) \leq \beta(G)\,\beta(F)$ for any two directed hypergraphs G and F.*

Proof. Let $\varepsilon > 0$ and, recalling the definition of β in (2.1) as the limit of β_t/t, let K be a sufficiently large integer such that for all $t \geq K$ we have $\beta_t(G)/t \leq \beta(G) + \varepsilon$ as well as $\beta_t(F)/t \leq \beta(F) + \varepsilon$. Let $\Sigma = \{0,1\}^K$ and consider the following scheme for the index coding problem on $G \bullet F$ with input alphabet Σ, which will consist of an inner and an outer code.

Let $\mathcal{E}_F$ denote an encoding function for F with input alphabet Σ achieving an optimal rate, i.e. minimizing $\log(|\Sigma_P|)/\log(|\Sigma|)$. For each $v \in V(G)$, the inner code applies $\mathcal{E}_F$ to the $|V(F)|$-tuple of messages indexed by the set $\{v\} \times V(F)$, obtaining a message m_v. Note that our assumption on $|\Sigma|$ implies that the length of m_v is equal to K' for some integer K' such that $K \leq K' \leq (\beta(F) + \varepsilon)K$. Next, let $\mathcal{E}_G$ denote an optimal encoding function for G with input $\{0,1\}^{K'}$. The outer code applies $\mathcal{E}_G$ to $\{m_v\}_{v \in V(G)}$ and the assumption on K ensures its output is at most $(\beta(G) + \varepsilon)K'$ bits long.

To verify that the scheme is a valid index code, consider a receiver in $G \bullet F$ represented by $e = ((w_G, w_F), (S_G \times V(F)) \cup (\{w_G\} \times S_F))$. To decode (w_G, w_F), the receiver first computes m_v for all $v \in S_G$. Since $\mathcal{E}_G$ is valid for G, receiver e can compute m_{w_G}, and since $\mathcal{E}_F$ is valid for F, this receiver can use the messages indexed by $\{w_G\} \times S_F$ along with m_{w_G} to compute (w_G, w_F).

Altogether, we have an encoding of K bits using at most $(\beta(F) + \varepsilon)(\beta(G) + \varepsilon)K$ bits of the public channel, and the required result follows from letting $\varepsilon \to 0$. ∎

[3]When there are n messages and exactly n receivers, w.l.o.g. receiver i wants the message x_i and one can encode the side-information by a graph on n vertices which contains the edge (i, j) iff receiver i knows the message x_j.

3 Linear programming

In this section we derive a linear program whose value constitutes a lower bound on the broadcast rate, and we prove that the value of the LP behaves supermultiplicatively under lexicographic products. In fact, rather than working with a specific linear program, we work with a general class of LP's having two types of constraints: those dictated by the network structure (which are the same for all LP's in the general class), and additional constraints depending only on the vertex set, generated by a *constraint schema*, i.e. a procedure for enumerating a finite set of constraints given an arbitrary finite index set. We identify some axioms on the constraint schema that constitute a sufficient condition for the LP value to be supermultiplicative. An example of a constraint schema which is important in network information theory is *submodularity*. For a given index set I, the submodularity schema enumerates all of the constraints of the form $z_S + z_T \geq z_{S\cap T} + z_{S\cup T}$ where S, T range over subsets of I.

Now we explain the general class of LPs which behave submultiplicatively under the lexicographic product and give bounds on β. Given an index code, if we sample each message independently and uniformly at random, we obtain a finite probability space on which the messages and the public channel are random variables. If S is a subset of these random variables, we will denote the Shannon entropy of the joint distribution of the variables in S by $H(S)$. If $S \subseteq T \subseteq \mathrm{cl}(S)$ then every message in $T \setminus S$ can be decoded given the messages in S and the public channel p, and consequently $H(S \cup \{p\}) = H(T \cup \{p\})$. More generally, if we normalize entropy (i.e. choose the base of the logarithm) so that $H(x) = 1$ for each message x, then for every $S \subseteq T$ we have

$$H(T \cup \{p\}) - H(S \cup \{p\}) \leq |T \setminus \mathrm{cl}(S)| \stackrel{\Delta}{=} c_{ST}, \tag{3.1}$$

where the above is taken as the definition of c_{ST}. This implies that for any index code we obtain a feasible solution of the primal LP in Figure 1 by setting $z_S = H(S \cup \{p\})$ for every S. Indeed, the first constraint expresses the fact that the value of p is determined by the values of the n messages, which are mutually independent. The second constraint was discussed above. The final line of the LP represents a set of constraints, corresponding to the rows of the matrix $A = (a_{qS})$, that are universally valid for any tuple of random variables indexed by the message set I. For instance, it is well known that the entropy of random variables has the submodularity property: $H(S) + H(T) \geq H(S \cup T) + H(S \cap T)$ if S, T are any two sets of random variables on the same sample space. So, for example, the rows of the constraint matrix A could be indexed by pairs of sets S, T, with entries in the (S, T) row chosen so that it represents the submodularity constraint (namely $a_{qS} = a_{qT} = 1$, $a_{q\,S\cap T} = a_{q\,S\cup T} = -1$ and all other entries of row a of A are zero). Noting that $H(\{p\}) \leq \beta(G)$ we can altogether conclude the following theorem.

Theorem 3.1. *For an index coding problem G, let $\mathfrak{B}(G)$ be the LP in Figure 1 when A represents the submodularity constraints and let $b(G)$ be its optimal solution. Then $b(G) \leq \beta(G)$.*

It is known that entropies of sets of random variables satisfy additional linear inequalities besides submodularity; if desired, the procedure for constructing the matrix A could be modified to incorporate some of these inequalities. Alternatively, in the context of restricted classes of encoding and decoding functions (e.g. linear functions) there may be additional inequalities that are specific to that class of functions, in which case the constraint matrix A may incorporate these inequalities and we obtain a linear program that is valid for this restricted model of index coding but not valid

min	$z_\emptyset$		max	$\lvert I\rvert \cdot w - \sum_{S\subset T} c_{ST} x_{ST}$	
s.t.	$z_I = \lvert I\rvert$	(w)	s.t.	$\sum_q a_{qS} y_q + \sum_{T\supset S} x_{ST} - \sum_{T\subset S} x_{TS} = 0$	$\forall S \neq \emptyset, I$
$\forall S \subset T$	$z_T - z_S \leq c_{ST}$	(x)		$\sum_q a_{q\emptyset} y_q + \sum_{T\neq\emptyset} x_{\emptyset T} = 1$	
	$Az \geq 0$	(y)		$\sum_q a_{qI} y_q - \sum_{T\neq I} x_{TI} + w = 0$	
				$x, y \geq 0$	

Figure 1: The LP and its dual.

in general. We will utilize such constraints in Section 6 when proving a separation between linear and non-linear network coding.

Definition 3.2. A *constraint schema* associates to each finite index set I a finite set $\mathcal{Q}(I)$ (indexing constraints) and a matrix $A(I)$ with rows indexed by $\mathcal{Q}(I)$ and columns indexed by $\mathcal{P}(I)$, the power set of I. In addition, to each Boolean lattice homomorphism[4] $h : \mathcal{P}(I) \to \mathcal{P}(J)$ it associates a function $h_* : \mathcal{Q}(I) \to \mathcal{Q}(J)$.

Let $\mathbf{1}$ be the $\mathcal{P}(I)$-indexed vector such that $\mathbf{1}_S = 1$ for all S, and let $\mathbf{1}_i$ be the vector where $(\mathbf{1}_i)_S = 1$ for all S containing i and otherwise $(\mathbf{1}_i)_S = 0$. We say that a constraint schema is *tight* if $A(I)\mathbf{1} = A(I)\mathbf{1}_i = 0$ for every index set I and element $i \in I$.

Given h and h_* let P_h and Q_h be matrices representing the linear transformations they induce on $\mathbb{R}^{\mathcal{P}(I)} \to \mathbb{R}^{\mathcal{P}(J)}$ and $\mathbb{R}^{\mathcal{Q}(I)} \to \mathbb{R}^{\mathcal{Q}(J)}$, respectively. That is, P_h and Q_h have zeros everywhere except $(P_h)_{h(S)S} = 1$ and $(Q_h)_{h_*(q)q} = 1$. We say that a constraint schema is *homomorphic* if it satisfies $A(J)^\mathsf{T} Q_h = P_h A(I)^\mathsf{T}$ for every Boolean lattice homomorphism $h : \mathcal{P}(I) \to \mathcal{P}(J)$.

Example 3.3. Earlier we alluded to the *submodularity constraint schema*. This is the constraint schema that associates to each index set I the constraint-index set $\mathcal{Q}(I) = \mathcal{P}(I) \times \mathcal{P}(I)$, along with the constraint matrix $A(I)$ whose entries are as follows. In row (S, T) and column U, we have an entry of 1 if $U = S$ or $U = T$, an entry of -1 if $U = S \cap T$ or $U = S \cup T$, and otherwise 0. (If any two of S, T, $S \cap T$, $S \cup T$ are equal, then that row of $A(I)$ is set to zero.) It is easy to verify that $A(I)\mathbf{1} = A(I)\mathbf{1}_i = 0$ for all $i \in I$, thus the schema is tight. For a homomorphism h, the corresponding mapping of constraint sets is $h_*(S, T) = (h(S), h(T))$. We claim that, equipped with this mapping of $h \to h_*$, the constraint schema is homomorphic. Indeed, to verify that $A(J)^\mathsf{T} Q_h = P_h A(I)^\mathsf{T}$ take any two sets $S, T \subset I$ and argue as follows to show that $u = P_h\, A(I)^\mathsf{T}\, e_{S,T}$ and $v = A(J)^\mathsf{T}\, Q_h\, e_{S,T}$ are identical (here and henceforth $e_{X,Y}$ denotes the standard basis vector of $\mathbb{R}^{\mathcal{P}(I)}$ having 1 in coordinate (X, Y) for $X, Y \subset I$). First observe that $A(I)^\mathsf{T}\, e_{S,T}$ is the vector $\tilde{u} \in \mathbb{R}^{\mathcal{P}(()I)}$ which has 0 entries everywhere except $\tilde{u}_S = \tilde{u}_T = 1$ and $\tilde{u}_{S\cup T} = \tilde{u}_{S\cap T} = -1$ provided that $S \not\subseteq T \not\subseteq S$, otherwise $\tilde{u} = 0$. As such, $u = P_h \tilde{u}$ has 0 entries everywhere except

$$u_{h(S)} = u_{h(T)} = 1\,, \quad u_{h(S\cup T)} = u_{h(S\cap T)} = -1$$

provided that $S \not\subseteq T \not\subseteq S$ and furthermore $h(S) \not\subseteq h(T) \not\subseteq h(S)$, otherwise $u = 0$ (for instance, if $S \subseteq T$ then $\tilde{u} = 0$ and so $u = 0$, whereas if $h(S) \subseteq h(T)$ then $\tilde{u}$ belongs to the kernel of P_h).

[4] A Boolean lattice homomorphism preserves unions and intersections, but does not necessarily map the empty set to the empty set nor the universal set to the universal set, and does not necessarily preserve complements.

Similarly, $Q_h\, e_{S,T} = e_{h(S),h(T)}$ and therefore $v = A(J)^\mathsf{T}\, e_{h(S),h(T)}$ has 0 entries everywhere except

$$v_{h(S)} = v_{h(T)} = 1\,, \quad v_{h(S)\cup h(T)} = v_{h(S)\cap h(T)} = -1$$

provided that $h(S) \not\subseteq h(T) \not\subseteq h(S)$, otherwise $v = 0$. To see that $u = v$ note that if $h(S) \subseteq h(T)$ then $u = v = 0$, and if $S \subseteq T$ then again we get $h(S) \subseteq h(T)$ due to monotonicity (recall that h is a lattice homomorphism) and so $u = v = 0$. Adding the analogous statements obtained from reversing the roles of S, T, it remains only to verify that $u = v$ in case $h(S) \not\subseteq h(T) \not\subseteq h(S)$, which reduces by the above definitions of u and v to requiring that $h(S\cup T) = h(S)\cup h(T)$ and $h(S\cap T) = h(S)\cap h(T)$. Both requirements are satisfied by definition of a Boolean lattice homomorphism, and altogether we conclude that the submodularity constraint schema is homomorphic.

Theorem 3.4. *Let A be a tight homomorphic constraint schema. For every index coding problem let $\rho(G)$ denote the optimum of the LP in Figure 1 when $I = V(G)$ and the constants c_{ST} are defined as in* (3.1)*. Then for every two index coding problems G and F, we have $\rho(G \bullet H) \geq \rho(G)\,\rho(F)$.*

Proof. It will be useful to rewrite the constraint set of the dual LP in a more succinct form. First, if x is any vector indexed by pairs S, T such that $S \subset T \subseteq I$, let $\nabla x \in \mathbb{R}^{\mathcal{P}(I)}$ denote the vector such that for all S, $(\nabla x)_S = \sum_{T\supset S} x_{ST} - \sum_{T\subset S} x_{TS}$. Next, for a set $S \subseteq I$, let e_S denote the standard basis vector vector in $\mathbb{R}^{\mathcal{P}(I)}$ whose S component is 1. Then the entire constraint set of the dual LP can be abbreviated to the following:

$$A^\mathsf{T} y + \nabla x + w e_I = e_\emptyset\,, \quad x, y \geq 0\,. \tag{3.2}$$

Some further simplifications of the dual can be obtained using the fact that the constraint schema is tight. For example, multiplying the left and right sides of (3.2) by the row vector $\mathbf{1}^\mathsf{T}$ gives

$$\mathbf{1}^\mathsf{T} A^\mathsf{T} y + \mathbf{1}^\mathsf{T}\nabla x + w = 1\,.$$

By the tightness of the constraint schema $\mathbf{1}^\mathsf{T} A^\mathsf{T} = 0$. It is straightforward to verify that $\mathbf{1}^\mathsf{T}\nabla x = 0$ and after eliminating these two terms from the equation above, we find simply that $w = 1$. Similarly, if we multiply the left and right sides of (3.2) by the row vector $\mathbf{1}_i^\mathsf{T}$ and substitute $w = 1$, we obtain $\mathbf{1}_i^\mathsf{T} A^\mathsf{T} y + \mathbf{1}_i^\mathsf{T}\nabla x + 1 = 0$ and consequently (again by the tightness) we arrive at $1 = -\mathbf{1}_i^\mathsf{T}\nabla x$. At the same time, $-\mathbf{1}_i^\mathsf{T}\nabla x = \sum_{\substack{S\subset T\\ i\in T\setminus S}} x_{ST}$ by definition of ∇x, hence summing over all $i \in I$ yields

$$|I| = \sum_{S\subset T} |T \setminus S|\, x_{ST}.$$

Plugging in this expression for $|I|$ and $w = 1$, the LP objective of the dual can be rewritten as

$$|I| - \sum_{S\subset T} c_{ST} x_{ST} = \sum_{S\subset T} \left(|T \setminus S| - c_{ST}\right)\, x_{ST} = \sum_{S\subset T} |T \cap (\mathrm{cl}(S) \setminus S)|\, x_{ST},$$

where the last equation used the fact that $c_{ST} = |T \setminus \mathrm{cl}(S)|$. We now define

$$d(S,T) = |T \cap (\mathrm{cl}(S) \setminus S)|$$

and altogether we arrive at the following reformulation of the dual LP.

$$\begin{aligned} \max \quad & \textstyle\sum_{S\subset T} d(S,T)\, x_{ST} \\ \text{s.t.} \quad & A^\mathsf{T} y + \nabla x = e_\emptyset - e_I \\ & x, y \geq 0\,. \end{aligned} \tag{3.3}$$

Now suppose that $(\xi^G, \eta^G), (\xi^F, \eta^F)$ are optimal solutions of the dual LP for G, F, achieving objective values $\rho(G)$ and $\rho(F)$, respectively. (Here ξ, η play the role of x, y from (3.3), resp.) We will show how to construct a pair of vectors $(\xi^{G\bullet F}, \eta^{G\bullet F})$ that is feasible for the dual LP of $G \bullet F$ and achieves an objective value of at least $\rho(G)\,\rho(F)$. The construction is as follows. Let $g : \mathcal{P}(V(G)) \to \mathcal{P}(V(G \bullet F))$ be the mapping $g(X) = X \times V(F)$. For sets $S \subset T \subseteq V(G)$, let $h^{ST} : \mathcal{P}(V(F)) \to \mathcal{P}(V(G \bullet F))$ be the mapping $h^{ST}(X) = (T \times X) \cup (S \times V(F))$. Observe that both mappings are Boolean lattice homomorphisms.

To gain intuition about the mappings g, h^{ST} it is useful to think of obtaining the vertex set of $G \bullet F$ by replacing every vertex of G with a copy of F. Here $g(\{v\})$ maps the vertex v in G to the copy of F that replaces v. The mapping $h^{ST}(\{u\})$ maps a vertex u in F to the vertex u in the copies of F that replace vertices in T, and then adds the set $\{u\} \times V(F)$.

Recall that Definition 3.2 associates two matrices P_h, Q_h to every Boolean lattice homomorphism $h : \mathcal{P}(I) \to \mathcal{P}(J)$. It is also useful to define a matrix R_h as follows: the columns and rows of R_h are indexed by pairs $S \subset T \subseteq I$ and $X \subset Y \subseteq J$, respectively, with the entry in row XY and column ST being equal to 1 if $X = h(S)$ and $Y = h(T)$, otherwise 0. Under this definition,

$$\nabla(R_h x) = P_h \nabla x \quad \text{for any } x \in \mathbb{R}^{\mathcal{P}(I)}. \tag{3.4}$$

Indeed, if $x = e_{S,T}$ for some $S \subset T \subseteq I$ then $\nabla e_{S,T} = e_S - e_T$ and so $P_h\, e_{S,T} = e_{h(S)} - e_{h(T)}$, whereas $\nabla(R_h e_{S,T}) = \nabla(e_{h(S),h(T)}) = e_{h(S)} - e_{h(T)}$.

We may now define

$$\xi^{G\bullet F} = \sum_{S\subset T} (\xi^G)_{ST}\, (R_{h^{ST}}\, \xi^F)\,, \tag{3.5}$$

$$\eta^{G\bullet F} = Q_g\, \eta^G + \sum_{S\subset T} (\xi^G)_{ST}\, (Q_{h^{ST}}\, \eta^F)\,. \tag{3.6}$$

In words, the dual solution for $G \bullet F$ contains a copy of the dual solution for F lifted according to h^{ST} for every pair $S \subset T$ and one copy of the dual solution of G lifted according to g. The feasibility of $(\xi^{G\bullet F}, \eta^{G\bullet F})$ will follow from multiple applications of the homomorphic property of the constraint schema and the feasibility of (ξ^F, η^F) and (ξ^G, η^G), achieved by the following claim.

Claim 3.5. *The pair $(\xi^{G\bullet F}, \eta^{G\bullet F})$ as defined in* (3.5),(3.6) *is a feasible dual solution.*

Proof. The matrices Q_g, $R_{h^{ST}}$, $Q_{h^{ST}}$ all have $\{0,1\}$-valued entries thus clearly $\xi^{G\bullet F}, \eta^{G\bullet F} \geq 0$. Letting $A = A(G \bullet F)$, we must prove that $A^{\mathsf{T}}\eta^{G\bullet F} + \nabla\xi^{G\bullet F} = e_\emptyset - e_{V(G\bullet F)}$. Plugging in the values of $(\xi^{G\bullet F}, \eta^{G\bullet F})$ we have

$$\begin{aligned} A^{\mathsf{T}}\eta^{G\bullet F} + \nabla\xi^{G\bullet F} &= A^{\mathsf{T}} Q_g \eta^G + \sum_{S\subset T} (\xi^G)_{ST}\, (A^{\mathsf{T}} Q_{h^{ST}} \eta^F) + \sum_{S\subset T} (\xi^G)_{ST}\, \nabla(R_{h^{ST}}\, \xi^F)\,, \\ &= P_g A(G)^{\mathsf{T}} \eta^G + \sum_{S\subset T} (\xi^G)_{ST}\, \left(P_{h^{ST}} A(F)^{\mathsf{T}} \eta^F + \nabla(R_{h^{ST}}\, \xi^F) \right). \end{aligned} \tag{3.7}$$

where the second equality applied the homomorphic property of the constraint schema. To treat the summation in the last expression above, recall (3.4) which implies that

$$P_{h^{ST}} A(F)^{\mathsf{T}} \eta^F + \nabla(R_{h^{ST}}\, \xi^F) = P_{h^{ST}} A(F)^{\mathsf{T}} \eta^F + P_{h^{ST}} \nabla\xi^F = P_{h^{ST}}(e_\emptyset - e_{V(F)})\,, \tag{3.8}$$

with the last equality due to the fact that (ξ^F, η^F) achieves the optimum of the dual LP for F. Recalling that $P_h e_S = e_{h(S)}$ for any h and combining it with the facts $h^{ST}(\emptyset) = S \times V(F)$ and $g(S) = S \times V(F)$ gives $P_{h^{ST}} e_\emptyset = e_{S \times V(F)} = P_g e_S$. Similarly, since $h^{ST}(V(F)) = T \times V(F)$ we have $P_{h^{ST}} e_{V(F)} = e_{T \times V(F)} = P_g e_T$, and plugging these identities in (3.8) combined with (3.7) gives:

$$A^\mathsf{T} \eta^{G \bullet F} + \nabla \xi^{G \bullet F} = P_g \Big[A(G)^\mathsf{T} \eta^G + \sum_{S \subset T} (\xi^G)_{ST} \, (e_S - e_T) \Big] .$$

Collecting together all the terms involving e_S for a given $S \in \mathcal{P}(I)$, we find that the coefficient of e_S is $\sum_{T \supset S} (\xi^G)_{ST} - \sum_{T \subset S} (\xi^G)_{ST} = (\nabla \xi^G)_S$. Hence,

$$A^\mathsf{T} \eta^{G \bullet F} + \nabla \xi^{G \bullet F} = P_g \left[A(G)^\mathsf{T} \eta^G + \nabla \xi^G \right] = P_g \left[e_\emptyset - e_{V(G)} \right] = e_\emptyset - e_{V(G \bullet F)} \, ,$$

where the second equality was due to (ξ^G, η^G) achieving the optimum of the dual LP for G. ■

To finish the proof, we must evaluate the dual LP objective and show that it is at least $\rho(G)\,\rho(F)$, as the next claim establishes:

Claim 3.6. *The LP objective for the dual solution given in Claim 3.5 has value at least $\rho(G)\,\rho(F)$.*

Proof. To simplify the notation, throughout this proof we will use K, L to denote subsets of $V(G \bullet F)$ while referring to subsets of $V(G)$ as S, T and to subsets of $V(F)$ as X, Y. We have

$$\begin{aligned} \sum_{K \subset L} d(K,L) (\xi^{G \bullet F})_{KL} &= \sum_{K \subset L} d(K,L) \sum_{S \subset T} (\xi^G)_{ST} \, (R_{h^{ST}} \, \xi^F)_{KL} \\ &= \sum_{S \subset T} (\xi^G)_{ST} \Big(\sum_{K \subset L} d(K,L) \, (R_{h^{ST}} \, \xi^F)_{KL} \Big) \\ &= \sum_{S \subset T} (\xi^G)_{ST} \Big(\sum_{X \subset Y} d\big(h^{ST}(X), h^{ST}(Y)\big) \, (\xi^F)_{XY} \Big) , \end{aligned} \tag{3.9}$$

where the last identity is by definition of R_h.

At this point we are interested in deriving a lower bound on $d\big(h^{ST}(X), h^{ST}(Y)\big)$, to which end we first need to analyze $\mathrm{cl}_{G \bullet F}(h^{ST}(X))$. Recall that $E(G \bullet F)$ consists of all hyperedges of the form (w, K) with $w = (w_G, w_F)$ and $K = (W_G \times V(F)) \cup (\{w_G\} \times W_F)$ for some pair of edges $(w_G, W_G) \in E(G)$ and $(w_F, W_F) \in E(F)$. We first claim that for any $S \subset T$ and $X \subset V(F)$,

$$\mathrm{cl}_{G \bullet F} \left(h^{ST}(X) \right) \setminus h^{ST}(X) \;\supseteq\; \Big(\big(\mathrm{cl}_G(S) \setminus S \big) \cap T \Big) \times \Big(\mathrm{cl}_F(X) \setminus X \Big) . \tag{3.10}$$

To show this, let $L \subseteq V(G \bullet F)$ denote the set on the right side of (3.10). Note that L contains no ordered pairs whose first component is in S or whose second component is in X, and therefore L is disjoint from $h^{ST}(X) = (T \times X) \cup (S \times V(F))$. Consequently, it suffices to show that $\mathrm{cl}_{G \bullet F} \left(h^{ST}(X) \right) \supseteq L$. Consider any $w = (w_G, w_F)$ belonging to L. As $w_G \in \mathrm{cl}_G(S) \setminus S$, there must exist an edge $(w_G, W_G) \in E(G)$ such that $W_G \subseteq S$. Similarly, there must exist an edge $(w_F, W_F) \in E(F)$ such that $W_F \subseteq X$. Recall from the definition of L that $\{w_G\} \subseteq T$. Now letting $K = (W_G \times V(F)) \cup (\{w_G\} \times W_F)$, we find that $K \subseteq (S \times V(F)) \cup (T \times X) = h^{ST}(X)$ and that $(w, K) \in E(G \bullet F)$, implying that $w \in \mathrm{cl}_{G \bullet F} \left(h^{ST}(X) \right)$ as desired.

Let $\hat{X} = h^{ST}(X)$ and $\hat{Y} = h^{ST}(Y)$, and recall that $d(\hat{X}, \hat{Y})$ is defined as $\left|\left(\mathrm{cl}_{G\bullet F}(\hat{X}) \setminus \hat{X}\right) \cap \hat{Y}\right|$. Using (3.10) and noting that $\hat{Y} \supseteq (T \times Y)$ we find that

$$\left(\mathrm{cl}_{G\bullet F}(\hat{X}) \setminus \hat{X}\right) \cap \hat{Y} \;\supseteq\; \Big(\left(\mathrm{cl}_G(S) \setminus S\right) \cap T\Big) \times \Big(\left(\mathrm{cl}_F(X) \setminus X\right) \cap Y\Big)$$

and hence

$$d(\hat{X}, \hat{Y}) \geq |(\mathrm{cl}_G(S) \setminus S) \cap T| \cdot |(\mathrm{cl}_F(X) \setminus X) \cap Y| = d(S,T)\, d(X,Y)\,.$$

Plugging this bound into (3.9) we find that

$$\begin{aligned}\sum_{K \subset L} d(K,L)(\xi^{G\bullet F})_{KL} &\geq \sum_{S \subset T} (\xi^G)_{ST} \sum_{X \subset Y} d(S,T) d(X,Y) (\xi^F)_{XY} \\ &= \Big(\sum_{S \subset T} d(S,T)(\xi^G)_{ST}\Big)\Big(\sum_{X \subset Y} d(X,Y)(\xi^F)_{XY}\Big) = \rho(G)\,\rho(F)\,,\end{aligned}$$

as required. ■

Combining Claims 3.5 and 3.6 concludes the proof of the Theorem 3.4. ■

Remark. The two sides of (3.10) are in fact equal for any non-degenerate index coding instances G and F, namely under the assumption that every $(w_G, W_G) \in E(G)$ has $w_G \notin W_G$ (otherwise this receiver already knows the required w_G and may be disregarded) and $W_G \neq \emptyset$ (otherwise the public channel must include w_G in plain form and we may disregard this message), and similarly for F. To see this, by definition of $\mathrm{cl}_{G\bullet F}(\cdot)$ and the fact that $h^{ST}(X) = (T \times X) \cup (S \times V(F))$ it suffices to show that every edge $(w, K) \in E(G \bullet F)$ with $K \subseteq h^{ST}(X)$ satisfies $w \in \left(\mathrm{cl}_G(S) \cap T\right) \times \mathrm{cl}_F(X)$. Take $(w, K) \in E(G \bullet F)$ and let $(w_G, W_G) \in E(G)$ and $(w_F, W_F) \in E(F)$ be the edges forming it as per Definition 2.1 of the lexicographic product. A prerequisite for $K \subseteq h^{ST}(X)$ is to have $w_G \in T$ as otherwise $\{w_G\} \times W_F \not\subseteq h^{ST}(X)$ (recall that $S \subset T$ and that $W_F \neq \emptyset$). Moreover, as X is strictly contained in $V(F)$ we must have $W_G \subseteq S$ in order to allow $W_G \times V(F) \subseteq h^{ST}(X)$, thus (using the fact that $w_G \notin W_G$ and so $w_G \notin S$) we further require that $W_F \subseteq X$. Altogether we have $W_G \subseteq S$, $W_F \subseteq X$ and $w_G \in T$, hence $(w_G, w_F) \in \left(\mathrm{cl}_G(S) \cap T\right) \times \mathrm{cl}_F(X)$ as required.

4 Separation between α and β

To prove Theorem 1.3, we start by using Theorem 3.1 to show that $\beta(C_5) > \alpha(C_5)$ where C_5 is the 5-cycle. Then we apply the power of Theorem 3.4 to transform this constant gap on C_5 to a polynomial gap on C_5^k.

First we show that $\beta(C_5) \geq b(C_5) \geq \frac{5}{2}$. We can show that $b(C_5) \geq \frac{5}{2}$ by providing a feasible dual solution for the LP $\mathfrak{B}$ with value $\frac{5}{2}$. This can easily be achieved by listing a set of primal constraints whose variables sum and cancel to show that $z_\emptyset \geq \frac{5}{2}$. Labeling the vertices of C_5 by $1, 2, 3, 4, 5$ sequentially, such a set of constraints is given below. It is helpful to note that in an index coding problem defined by an undirected graph, $x \in \mathrm{cl}(S)$ if $x \in S$ or all the neighbors of x

are in S.

$$\begin{aligned}
2 &\geq z_{\{1,3\}} - z_{\emptyset} \\
2 &\geq z_{\{2,4\}} - z_{\emptyset} \\
1 &\geq z_{\{5\}} - z_{\emptyset} \\
0 &\geq z_{\{1,2,3\}} - z_{\{1,3\}} \\
0 &\geq z_{\{2,3,4\}} - z_{\{2,4\}} \\
z_{\{2,3,4\}} + z_{\{1,2,3\}} &\geq z_{\{2,3\}} + z_{\{1,2,3,4\}} \\
z_{\{2,3\}} + z_{\{5\}} &\geq z_{\emptyset} + z_{\{2,3,5\}} \\
0 &\geq z_{\{1,2,3,4,5\}} - z_{\{1,2,3,4\}} \\
0 &\geq z_{\{1,2,3,4,5\}} - z_{\{2,3,5\}} \\
z_{\{1,2,3,4,5\}} &= 5 \\
z_{\{1,2,3,4,5\}} &= 5
\end{aligned}$$

Applying Theorem 3.4 we deduce that for any integer $k \geq 1$ the k-th lexicographic power of C_5 satisfies $\beta(C_5^k) \geq b(C_5^k) \geq \left(\frac{5}{2}\right)^k$. Furthermore, $\alpha(C_5) = 2$ and it is well known that the independence number is multiplicative on lexicographic products and so $\alpha(C_5^k) = 2^k$. Altogether, C_5^k is a graph on $n = 5^k$ vertices with $\alpha = n^{\log_5(2)}$ and $\beta \geq n^{1-\log_5(2)}$, implying our result.

5 Matroids and index coding

Recall that a matroid is a pair $M = (E, r)$ where E is a ground set and $r : 2^E \to \mathbb{N}$ is a rank function satisfying

(i) $r(A) \leq |A|$ for all $A \subseteq E$;
(ii) $r(A) \leq r(B)$ for all $A \subseteq B \subseteq E$ (monotonicity);
(iii) $r(A) + r(B) \geq r(A \cup B) + r(A \cap B)$ for all $A, B \subseteq E$ (submodularity).

The rank vector of a matroid, $\vec{\mathbf{r}}(M)$, is a $2^{|E|}$-dimensional vector indexed by subsets of E, such that its S-th coordinate is $r(S)$. A subset $S \subseteq E$ is called *independent* if $r(S) = |S|$ and it is called a *basis* of M if $r(S) = |S| = r(E)$.

In this section we give a construction mapping a matroid to an instance of index coding that exactly captures the dependencies in the matroid. We proceed to show some useful connections between matroid properties and the broadcast rate of the corresponding index coding problem.

Definition 5.1. Let $M = (E, r)$ be a matroid. The hypergraph index coding problem *associated* to M, denoted by G_M, has a message set E and all receivers of the form

$$\left\{(x, S) \mid x \in E\,,\, S \subseteq E\,,\, r(S) = r(S \cup \{x\})\right\}.$$

Remark. A similar yet slightly more complicated construction was given in [10]. Our construction is (essentially) a subset of the one appearing there. A construction that maps a matroid to a network coding problem is given in [6, 7]. They prove an analog of Proposition 5.2.

Proposition 5.2. *For a matroid $M = (E, r)$, $b(G_M) = |E| - r(E)$.*

Proof. In what follows we will let $n = |E|$ and $r = r(E)$. To show that $b(G_M) \leq n - r$ it suffices to show $z_S = r(S) + n - r$ is a feasible primal solution to the LP $\mathfrak{B}(G_M)$. The feasibility of constraints

(w) and (x) follows trivially from the definition of G_M and properties of a matroid. The feasibility of $(y) : z_T - z_S \leq c_{ST}\ \forall S \subset T$ follows from repeated application of submodularity:

$$\begin{aligned} z_T - z_S = r(T) - r(S) &\leq \sum_{x \in T \setminus S} r(S \cup \{x\}) - r(S) \\ &\leq \sum_{x \in \mathrm{cl}(S)} (r(S \cup \{x\}) - r(S)) \ + \sum_{x \in T \setminus \mathrm{cl}(S)} r(\{x\}) \leq |T \setminus \mathrm{cl}(S)| = c_{ST}. \end{aligned}$$

To prove the reverse inequality, let S be any basis of M and note that $z_\emptyset = z_E - (z_E - z_S) - (z_S - z_\emptyset) \geq n - c_{SE} - c_{\emptyset S} = n - r$. ■

The following definition relaxes the notion of a representation for a matroid.

Definition 5.3. A matroid $M = (E, r)$ with $|E| = n$ is *under-representable* in d dimensions over a finite field $\mathbb{F}$ if there exists a $d \times n$ matrix with entries in $\mathbb{F}$ and columns indexed by elements of E such that (i) the rows are independent and (ii) if $r(x \cup S) = r(S)$ then the columns indexed by $x \cup S$ are dependent.

Observe that if a matrix represents M then it also under-represents M. We next show a relation between under-representations for M over $\mathbb{F}$ and the *scalar* linear rate $\lambda_1^{\mathbb{F}}$, where the alphabet vector space, over which the encoding functions are required to be linear, is single-dimensional. Note that $\lambda^{\mathbb{F}} \leq \lambda_1^{\mathbb{F}}$. The following is the analogue of Theorem 8 in [10] for our version of the matroid to index coding mapping.

Theorem 5.4. *A matroid $M = (E, r)$ with $|E| = n$ is under-representable in d dimensions over a finite field $\mathbb{F}$ if and only if $\lambda_1^{\mathbb{F}}(G_M) \leq n - d$. In particular, if M is representable over $\mathbb{F}$ then $\lambda^{\mathbb{F}}(G_M) = \beta(G_M) = n - r(E)$.*

Proof. Let R be a $d \times n$ matrix which under-represents M in d dimensions over $\mathbb{F}$. Let Q be an $(n-d) \times n$ matrix whose rows span the kernel of R. We will show that Q is a valid encoding matrix for G_M. Let $y \in \mathbb{F}^E$ be some input message set and consider a receiver (x, S), who wishes to decode y_x from $\{y_z : z \in S\}$ and the broadcast message Qy. Extend $\ker(Q)$ arbitrarily into a basis B for $\mathbb{F}^E$ and let $y = y' + y''$ be the unique decomposition according to B such that $y' \in \ker(Q)$. Clearly, $Qy'' = Qy$ since $y' \in \ker(Q)$, hence one can recover y'' from the public channel by triangulating Q. It remains for the receiver (x, S) to recover y'_x. To this end, observe that the rows of R span $\ker(Q)$ and recall that by Definitions 5.1 and 5.3, column x of R is a linear combination of the columns of R indexed by S. Since y' is in the row-space of R it follows that y'_x is equal to the exact same linear combination of the components of y' indexed by S, all of which are known to the receiver. Altogether, the receiver can recover both y'_x and y''_x and obtain the message x. As this holds for any receiver, we conclude that Q is a valid encoding matrix and thus $\lambda_1^{\mathbb{F}}(G_M) \leq n - d$. When $d = r(E)$ the inequality is tight because this upper bound coincides with the lower bound given by Proposition 5.2.

Conversely, suppose that there exists a scalar linear code for G_M over $\mathbb{F}$ with rate $n - d$, and let Q be a corresponding $(n-d) \times n$ encoding matrix of rank $n - d$. Let R be a $d \times n$ matrix whose rows span the kernel of Q. We claim that R under-represents M. Indeed, consider a receiver (x, S).

It is easy to verify that this receiver has a linear decoding function[5] of the form $u^{\mathsf{T}} \cdot Qy + v^{\mathsf{T}} \cdot y_S$ for some vectors u, v, where y_S is the vector formed by restricting y to the indices of S. As Q is a valid encoding matrix for G_M, this evaluates to y_x for any $y \in \mathbb{F}^E$. In particular, if y^{T} is a row of R then $Qy = 0$ and so $v^{\mathsf{T}} \cdot y_S = y_x$, and applying this argument to every row of R verifies that column x of R is a linear combination of the columns of R indexed by S (with coefficients from v). Since this holds for any receiver we have that R under-represents M, as required. ■

We conclude this section with a result that will be useful in establishing lower bounds on the value of the LP for G_M with a given constraint matrix A.

Theorem 5.5. *Suppose that $M = (E, r)$ is a matroid and A is a matrix such that $A\mathbf{1} = 0$ and $A\vec{\mathbf{r}}(M) \not\geq 0$. If the linear program in Figure 1 is instantiated with constraint matrix A, then the value of the LP is strictly greater than $|E| - r(E)$.*

Proof. We will give a dual solution (w, x, y) to the LP with value strictly greater than $|E| - r(E)$.

Recalling the hypothesis $A\vec{\mathbf{r}}(M) \not\geq 0$, let q be a row of A such that $\sum_{S \subseteq E} a_{qS} r(S) < 0$. Let $\mathcal{S}^+ = \{S \subseteq E \mid a_{qS} > 0, S \neq E, \emptyset\}$ and $\mathcal{S}^- = \{S \subseteq E \mid a_{qS} < 0, S \neq E, \emptyset\}$. Note that the hypothesis that $A\mathbf{1} = 0$ implies that $a_{q\emptyset} + \sum_{S \in \mathcal{S}^+} a_{qS} = -\left(a_{qE} + \sum_{S \in \mathcal{S}^-} a_{qS}\right)$. Assume that A is scaled so $a_{q\emptyset} + \sum_{S \in \mathcal{S}^+} a_{qS} = -\left(a_{qE} + \sum_{S \in \mathcal{S}^-} a_{qS}\right) = 1$. This assumption is without loss of generality since $a_{qE} + \sum_{S \in \mathcal{S}^-} a_{qS}$ is strictly negative, as can be seen from the following calculation:

$$\begin{aligned} r(E)\left(a_{qE} + \sum_{S \in \mathcal{S}^-} a_{qS}\right) &\leq a_{qE} r(E) + \sum_{S \in \mathcal{S}^-} a_{qS} r(S) \leq a_{qE} r(E) + \sum_{S \in \mathcal{S}^-} a_{qS} r(S) + \sum_{S \in \mathcal{S}^+} a_{qS} r(S) \\ &= \sum_{S} a_{qS} r(S) \;\; < \;\; 0\,. \end{aligned}$$

Define the dual vector y by setting $y_q = 1$ and $y_{q'} = 0$ for rows $q' \neq q$ of A. To define the dual vector x, let us first associate to every set $S \subseteq E$ a matroid basis $b(S)$ such that the set $m(S) = b(S) \cap S$ is a maximal independent subset of S, i.e. $|m(S)| = r(m(S)) = r(S)$. Let $u(S) = S \cup b(S)$. For every $S \in \mathcal{S}^+$, let $x_{\emptyset m(S)} = x_{m(S)S} = a_{qS}$ and for every $S \in \mathcal{S}^-$, let $x_{Su(S)} = x_{u(S)E} = -a_{qS}$. Set all other values of x_{ST} to zero. Finally, set $w = 1$. By construction, (w, x, y) satisfies all of the dual constraints. Using the relations $c_{\emptyset m(S)} = r(S)$, $c_{Su(S)} = r(E) - r(S)$, $c_{m(S)S} = c_{u(S)E} = 0$, we find that the dual LP objective value is

$$\begin{aligned} |E|\, w - \sum_{S \subset T} c_{ST} x_{ST} &= |E| - \sum_{S \in \mathcal{S}^+} (c_{\emptyset m(S)} + c_{m(S)S}) a_{qS} - \sum_{S \in \mathcal{S}^-} (c_{Su(S)} + c_{u(S)E})(-a_{qS}) \\ &= |E| - \sum_{S \in \mathcal{S}^+} r(S) a_{qS} + \sum_{S \in \mathcal{S}^-} (r(E) - r(S)) a_{qS} \\ &= |E| + \sum_{S \in \mathcal{S}^-} a_{qS} r(E) - \sum_{S} a_{qS} r(S) + a_{q\emptyset} r(\emptyset) + a_{qE} r(E) \\ &= |E| - r(E) - \sum_{S} a_{qS} r(S). \end{aligned}$$

By hypothesis $\sum_S a_{qS} r(S) < 0$, and the proposition follows. ■

[5] This follows e.g. from decomposing y as above into $y' + y''$ where $y' \in \ker(Q)$. By definition y''_x is a linear combination of the Qy entries. Similarly, y'_x must be a linear combination of $\{y_z : z \in S\}$, otherwise there would exist some $y \in \ker(Q)$ with $y_x \neq 0$ and $y_z = 0$ for all $z \in S$, making it indistinguishable to this receiver from $y = 0$.

6 Separation between linear and non-linear rates

In this section we prove Theorem 1.2. To this end we will first show that the linear rate over a field of even characteristic is strictly better than the linear rate over a field of odd characteristic for the index coding problem associated to the Fano matroid, and that the reverse relation holds for the non-Fano matroid. Then we will take the lexicographic product of the two index codes to get a gap between the linear and non-linear coding rates, and then use lexicographic products again to amplify that gap.

The *Fano matroid*, denoted $\mathcal{F}$, and the *non-Fano matroid*, denoted $\mathcal{N}$, are 7 element, rank 3 matroids. The seven columns of the matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 1 \end{pmatrix}$$

constitute a linear representation of the Fano matroid when $\text{char}(\mathbb{F}) = 2$ and one for the non-Fano matroid when $\text{char}(\mathbb{F}) \neq 2$. We will use $\mathcal{U} = \{100, 010, 001, 110, 101, 011, 111\}$ to index the elements of the two matroids. Further let $\mathcal{O} \subset \mathcal{U}$ be the vectors with odd hamming weight, let $\mathcal{B}$ be the vectors with hamming weight one and let $i + j$ for $i, j \in \mathcal{U}$ be the bitwise addition of i, j.

It is well known that the Fano matroid is representable only in a field of characteristic 2, and the non-Fano matroid is representable in any field whose characteristic is different from 2 but not in fields of characteristic 2. We use a generalization of this fact to prove the following theorem that directly implies Theorem 1.2.

Theorem 6.1 (Separation Theorem)**.** *Let $G = G_{\mathcal{F}} \bullet G_{\mathcal{N}}$. There exists some $\varepsilon > 0$ such that $\beta(G^{\bullet n}) = 16^n$ whereas $\lambda(G^{\bullet n}) \geq (16 + \varepsilon)^n$ for all n.*

The fact that $\beta(G^{\bullet n}) = 16^n$ will be a straightforward application of Proposition 5.2 and Theorem 5.4. The lower bound on the linear rate however will require considerably more effort. In order to bound λ from below we will extend the LP $\mathfrak{B}$ to two LPs, one of which will be a lower bound for linear codes over fields with odd characteristic and the other for linear codes over even characteristic. Each one will supplement the matrix A in the LP with a set of constraints, one set derived from dimension inequalities based on the representation of the Fano matroid and the other from the non-Fano matroid. The LP that gives a lower bound for linear codes over a field with even characteristic will be used to show that the linear broadcast rate of $G_{\mathcal{N}}$ over a field of even characteristic is strictly greater than four, and the LP for odd characteristic will imply the corresponding result for $G_{\mathcal{F}}$. Furthermore, the constraints will satisfy the conditions of Theorem 3.4. Putting this all together implies that when we take the lexicographic product of the Fano and non-Fano index coding problems, no linear code is as good as one that combines linear codes over $\mathbb{F}_2$ and $\mathbb{F}_3$.

Before explaining how we derive these constraints, we introduce a bit of notation. If $\{V_i\}_{i \in I}$ are subspaces of a vector space V, let the span of V_i and V_j be denoted $V_i + V_j$ and let $\dim(\{V_i\}_{i \in I})$ be the dimension of the span of $\{V_i\}_{i \in I}$. Also, let $\vec{\mathbf{d}}(\{V_i\}_{i \in I})$ be a $2^{|I|}$ dimensional vector indexed by the subsets of I such that the coordinate indexed by S is $\dim(\{V_i\}_{i \in S})$. We let $V_1 \oplus \cdots \oplus V_k$ denote the sum of mutually complementary subspaces $V_1, \ldots, V_k$. If $V = V_1 \oplus \cdots \oplus V_k$ then V is isomorphic to the vector space $\prod_{i=1}^{k} V_i$ via the mapping $(v_1, \ldots, v_k) \mapsto v_1 + \cdots + v_k$. In this case, for an index set $S \subseteq \{1, \ldots, k\}$, we will use π_S to denote the projection function $V \to \oplus_{i \in S} V_i$, i.e. the function that maps an element $v = \sum_{i=1}^{k} v_i$ to the element $\pi_S(v) = \sum_{i \in S} v_i$.

The fact that the Fano matroid can be represented over $\mathbb{F}_2$ and the non-Fano matroid cannot tells us something about dimension dependencies that can occur in $\mathbb{F}_2$. The following lemma is extracting the critical dimension relations that distinguish vector spaces over $\mathbb{F}$ with $\text{char}(\mathbb{F}) = 2$.

Lemma 6.2. *Let $V = V_1 \oplus V_2 \oplus V_3$ be a vector space over a field $\mathbb{F}$, and suppose $W \subset V$ is a linear subspace that is complementary to each of $V_1 \oplus V_2, V_1 \oplus V_3, V_2 \oplus V_3$. Then*

$$\dim(\pi_{12}(W), \pi_{13}(W), \pi_{23}(W)) = \begin{cases} 2\dim(W) & \text{if } \text{char}(\mathbb{F}) = 2 \\ 3\dim(W) & \text{if } \text{char}(\mathbb{F}) \neq 2. \end{cases} \tag{6.1}$$

Proof. Recalling that V is isomorphic to $\prod_{i=1}^{3} V_i$, we will write elements of V as ordered triples. Our assumption that W is complementary to each of $V_1 \oplus V_2, V_1 \oplus V_3, V_2 \oplus V_3$ implies that a nonzero element of W has three nonzero coordinates, a fact that we will use in both cases of the lemma.

If $\text{char}(\mathbb{F}) = 2$, then every vector $(x, y, z) \in V$ satisfies

$$\pi_{12}(x, y, z) + \pi_{13}(x, y, z) = (x, y, 0) + (x, 0, z) = (0, y, z) = \pi_{23}(x, y, z)$$

hence $\pi_{12}(W) + \pi_{13}(W) = \pi_{23}(W)$. Consequently

$$\dim(\pi_{12}(W), \pi_{13}(W), \pi_{23}(W)) = \dim(\pi_{12}(W), \pi_{13}(W)) \leq 2\dim(W).$$

To prove the reverse inequality we observe that $\pi_{12}(W)$ and $\pi_{13}(W)$ are complementary, since every nonzero element of $\pi_{12}(W)$ is of the form $(x, y, 0)$ with $x, y \neq 0$, whereas every nonzero element of $\pi_{13}(W)$ is of the form $(x, 0, z)$ with $x, z \neq 0$, and hence $\pi_{12}(W) \cap \pi_{13}(W) = \{0\}$.

When $\text{char}(\mathbb{F}) \neq 2$, we prove Equation (6.1) by showing that $\pi_{12}(W), \pi_{13}(W), \pi_{23}(W)$ are mutually complementary. Consider any three vectors $w_1 = (x_1, y_1, z_1)$, $w_2 = (x_2, y_2, z_2)$, and $w_3 = (x_3, y_3, z_3)$, all belonging to W, such that

$$0 = \pi_{23}(x_1, y_1, z_1) + \pi_{13}(x_2, y_2, z_2) + \pi_{12}(x_3, y_3, z_3) = (x_2 + x_3, y_1 + y_3, z_1 + z_2)\,.$$

This implies that $x_2 + x_3 = 0$, so the first coordinate of $w_2 + w_3$ is zero. However, the zero vector is the only vector in W whose first coordinate is zero, hence $w_2 + w_3 = 0$. Similarly, $w_1 + w_3 = 0$ and $w_1 + w_2 = 0$. Now using the fact that 2 is invertible in $\mathbb{F}$, we deduce that $w_1 = \frac{1}{2}[(w_1 + w_2) + (w_1 + w_3) - (w_2 + w_3)] = 0$, and similarly $w_2 = 0$ and $w_3 = 0$. Thus, the only way to express the zero vector as a sum of vectors in $\pi_{12}(W), \pi_{13}(W), \pi_{23}(W)$ is if all three summands are zero, i.e. those three subspaces are mutually complementary as claimed. ∎

6.1 Linear codes over fields of characteristic two

This section provides the ingredients for proving that $\lambda^{\mathbb{F}}(G_{\mathcal{F}}) > 4$ for $\mathbb{F}$ with $\text{char}(\mathbb{F}) = 2$.

Lemma 6.3 (Conditional Even Characteristic Inequality)**.** *Suppose $\{V_i\}_{i \in \mathcal{U}}$ are 7 subspaces of a vector space over $\mathbb{F}$ such that* $\text{char}(\mathbb{F}) = 2$ *and*

(i) $\dim(\{V_i\}_{i \in \mathcal{O}}) = \dim(\{V_i\}_{i \in \mathcal{B}})$

(ii) $\dim(V_i, V_j, V_k) = \dim(V_i) + \dim(V_j) + \dim(V_k)\ \forall i, j, k \in \mathcal{O}$

(iii) $\dim(V_i, V_j, V_{i+j}) = \dim(V_i, V_j)\ \forall i, j \in \mathcal{O}$

Then $\dim(V_{110}, V_{101}, V_{011}) \leq 2\dim(V_{111})$.

Proof of Lemma 6.3. Hypotheses (i) and (iii) of the lemma imply that all 7 subspaces are contained in the span of $V_{100}, V_{010}, V_{001}$. Moreover, hypothesis (ii) implies that $V_{100}, V_{010}, V_{001}$ are mutually complementary and that V_{111} is complementary to each of $V_{100} + V_{010}$, $V_{100} + V_{001}$, $V_{010} + V_{001}$. Thus, we can apply Lemma 6.2 with $V = V_{100} \oplus V_{010} \oplus V_{001}$ and $W = V_{111}$, yielding the equation $\dim(\pi_{12}(V_{111}), \pi_{23}(V_{111}), \pi_{13}(V_{111})) = 2\dim(V_{111})$.

We claim that $\pi_{12}(V_{111}) = (V_{001} + V_{111}) \cap (V_{100} + V_{010})$. To see this, take an arbitrary element $w \in V_{111}$ having a unique representation of the form $x + y + z$ with $x \in V_{100}, y \in V_{010}, z \in V_{001}$. By definition $\pi_{12}(w) = x + y = w - z$, from which it can be seen at once that $\pi_{12}(w)$ belongs to both $V_{100} + V_{010}$ and $V_{001} + V_{111}$. Conversely, any element $v \in (V_{001} + V_{111}) \cap (V_{100} + V_{010})$ can be expressed as $v = w - z$ where $w \in V_{111}, z \in V_{001}$ but it can also be expressed as $v = x + y$ where $x \in V_{100}, y \in V_{010}$. Consequently, $w = x + y + z$ and $v = \pi_{12}(w)$.

Hypothesis (iii) implies that V_{110} is contained in both $V_{001} + V_{111}$ and $V_{100} + V_{010}$, hence $V_{110} \subseteq \pi_{12}(V_{111})$. Similarly $V_{101} \subseteq \pi_{13}(V_{111})$ and $V_{011} \subseteq \pi_{23}(V_{111})$. Hence $\dim(V_{110}, V_{101}, V_{011}) \leq \dim(\pi_{12}(V_{111}), \pi_{23}(V_{111}), \pi_{13}(V_{111})) = 2\dim(V_{111})$, as desired. ■

In what follows we will transform the conditional inequalities given in the lemma above to a general inequality that applies to any 7 subspaces of a vector space over a field of characteristic 2 by using the following approach. We will start with arbitrary subspaces and then repeatedly modify them until they satisfy the conditions of Lemma 6.3. At that point the result in this conditional lemma will imply an inequality involving the dimensions of the modified subspaces, which we will express in terms of the dimensions of the original subspaces.

Theorem 6.4 (Even Characteristic Inequality)**.** *There exists a* 2^7*-dimensional vector* Λ_{even} *such that for any 7 subspaces* $\{V_i\}_{i\in\mathcal{U}}$ *of a vector space over* $\mathbb{F}$ *with* $\text{char}(\mathbb{F}) = 2$*,*

$$\Lambda_{\text{even}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i\in\mathcal{U}}) \geq 0 \text{ and } \Lambda_{\text{even}} \cdot \vec{\mathbf{r}}(\mathcal{N}) < 0.$$

Proof. As mentioned above, the proof will proceed by repeatedly modifying the input subspaces until they satisfy the requirements of Lemma 6.3. The modifications we make to a vector space are of one type: we *delete* a vector w from a subspace V that contains w, by letting B be a basis of V containing w and then replacing V with the span of $B \setminus w$.

Let $\{V_i\}_{i\in\mathcal{U}}$ be seven subspaces of a vector space V over $\mathbb{F}$ such that $\text{char}(\mathbb{F}) = 2$. We will modify the subspaces $\{V_i\}_{i\in\mathcal{U}}$ into $\{V_i'\}_{i\in\mathcal{U}}$ that satisfy the conditions of Lemma 6.3. To start, we set $\{V_i'\}_{i\in\mathcal{U}} = \{V_i\}_{i\in\mathcal{U}}$. We then update $\{V_i'\}_{i\in\mathcal{U}}$ in three steps, each of which deletes vectors of a certain type in an iterative fashion. The order of the deletions within each step is arbitrary.

Step 1: Vectors in V_{111}' but not in $\sum_{i\in\mathcal{B}} V_i'$ from V_{111}'.

Step 2: (a) Vectors in $V_{100}' \cap V_{010}'$ from V_{010}'.
(b) Vectors in $V_{001}' \cap (V_{100}' + V_{010}')$ from V_{001}'.
(c) Vectors in $V_{111}' \cap (V_{100}' + V_{010}')$ from V_{111}'.
(d) Vectors in $V_{111}' \cap (V_{010}' + V_{001}')$ from V_{111}'.
(e) Vectors in $V_{111}' \cap (V_{100}' + V_{001}')$ from V_{111}'.

Step 3: Vectors in V'_{i+j} but not in $V'_i + V'_j$ for $i, j \in \mathcal{O}$ from V'_{i+j}.

First, we argue that $\{V'_i\}_{i\in\mathcal{U}}$ satisfy the conditions of Lemma 6.3. The deletions in step (1) ensure that V'_{111} is contained in $\sum_{i\in\mathcal{B}} V'_i$, thus satisfying condition (i). The deletions in steps (2a)–(2b) ensure that $V'_{100}, V'_{010}, V'_{001}$ are mutually complementary, and steps (2c)–(2d) ensure that V'_{111} is complementary to the sum of any two of them, thus satisfying condition (ii). Furthermore, step (2) does not change $\sum_{i\in\mathcal{B}} V'_i$ because we only delete a vector from one of $\{V'_i\}_{i\in\mathcal{B}}$ when it belongs to the span of the other two. Thus condition (i) is still satisfied at the end of step (2). Step (3) ensures that V'_{i+j} is contained in $V'_i + V'_j$, thus satisfying condition (iii). Furthermore, it does not modify $V'_i, i \in \mathcal{O}$, and thus conditions (i) and (ii) remain satisfied after step (3).

Now, by Lemma 6.3 we have that

$$\dim(V'_{110}, V'_{101}, V'_{011}) \le 2\dim(V'_{111}). \tag{6.2}$$

Let

$$\begin{aligned}
\delta &= \dim(V_{111}, \{V_i\}_{i\in\mathcal{B}}) - \dim(\{V_i\}_{i\in\mathcal{B}}) \\
\delta[i|j,k] &= \dim(V_i, V_j, V_k) - \dim(V_j, V_k) \\
\delta[i;j] &= \dim(V_i \cap V_j) = \dim(V_i) + \dim(V_j) - \dim(V_i, V_j) \\
\delta[i;j,k] &= \dim(V_i \cap (V_j + V_k)) = \dim(V_i) + \dim(V_j, V_k) - \dim(V_i, V_j, V_k)
\end{aligned}$$

Observe that after step (1) $\dim(V'_{111}) = \dim(V_{111}) - \delta$, and steps (2) and (3) only delete more vectors from V'_{111}, so we have $\dim(V'_{111}) \le \dim(V_{111}) - \delta$.

It remains to get a lower bound on $\dim(V'_{110}, V'_{101}, V'_{011})$ in terms of dimensions of subsets of $\{V_i\}_{i\in\mathcal{U}}$. We do this by giving an upper bound on the total number of vectors deleted from $E = V'_{110}+V'_{101}+V'_{011}$ in terms of the δ terms we defined above. In steps (1) and (2) we delete nothing from E, but we delete some vectors from $V'_i, i \in \mathcal{O}$. Specifically, $\delta[100;010]$ vectors are deleted from V'_{010}, $\delta[001;100,010]$ vectors are deleted from V'_{001}, and no vectors are deleted from V_{100}. As already noted, step (1) deletes δ vectors from V'_{111}, while step (2) deletes at most $\sum_{i,j\in\mathcal{B}} \delta[111;i,j]$ vectors from V'_{111}. To summarize, the dimensions of $V'_i, i \in \mathcal{O}$, after steps (1) and (2), satisfy:

$$\dim(V'_{100}) = \dim(V_{100}) \tag{6.3}$$

$$\dim(V'_{010}) = \dim(V_{010}) - \delta[100;010] \tag{6.4}$$

$$\dim(V'_{001}) = \dim(V_{001}) - \delta[001;100,010] \tag{6.5}$$

$$\dim(V'_{111}) \ge \dim(V_{111}) - \delta - \sum_{i,j\in\mathcal{B}} \delta[111;i,j]. \tag{6.6}$$

In step (3), when we delete vectors in V'_{i+j} but not in $V'_i + V'_j$; if no deletions had taken place in prior steps then the number of vectors deleted from V'_{i+j} would be $\delta[i+j|i,j]$. However, the deletions that took place in steps (1) and (2) have the effect of reducing the dimension of $V'_i + V'_j$, and we must adjust our upper bound on the number of vectors deleted from V'_{i+j} to account for the potential difference in dimension between $V_i + V_j$ and $V'_i + V'_j$. When $i = 100$, $j = 010$, there is no difference between $V_i + V_j$ and $V'_i + V'_j$, because the only time vectors are deleted from either one of these subspaces is in step (2a), when vectors in $V'_{100} \cap V'_{010}$ are deleted from V'_{010} without changing the dimension of $V'_{100} + V'_{010}$. For all other pairs $i, j \in \mathcal{O}$, we use the upper bound

$$\dim(V_i + V_j) - \dim(V'_i + V'_j) \le \left[\dim(V_i) - \dim(V'_i)\right] + \left[\dim(V_j) - \dim(V'_j)\right],$$

which is valid for any four subspaces V_i, V_j, V_i', V_j' satisfying $V_i' \subseteq V_i$, $V_j' \subseteq V_j$. Let $\Delta\dim(V_i)$ denote the difference $\dim(V_i) - \dim(V_i')$. Combining these upper bounds, we find that the number of extra vectors deleted from E in step (3) because of differences in dimension between $V_i' + V_j'$ and $V_i + V_j$ is at most

$$\begin{aligned}\left(\sum_{i,j\in\mathcal{O}} \Delta\dim(V_i) + \Delta\dim(V_j)\right) &- \Delta\dim(V_{100}) - \Delta\dim(V_{010}) \\ &= 2\left(\sum_{i\in\{100,010\}} \Delta\dim(V_i)\right) + 3\left(\sum_{i\in\{001,111\}} \Delta\dim(V_i)\right) \\ &\leq 2\delta[100;010] + 3\delta[001;100,010] + 3\delta + 3\sum_{i,j\in\mathcal{B}} \delta[111;i,j]\end{aligned}$$

where the last inequality follows by combining equations (6.3)–(6.6).

We now sum up our upper bounds on the number of vectors deleted from E in step (3), to find that

$$\dim(E) \geq \dim(V_{110}, V_{101}, V_{011}) - \sum_{i,j\in\mathcal{O}} \delta[i{+}j|i,j] - 2\delta[100;010] - 3\delta[001;100,010] - 3\delta - 3\sum_{i,j\in\mathcal{B}} \delta[111;i,j]. \tag{6.7}$$

Expanding out all the δ terms, combining with the upper bound $\dim(V_{111}') \leq \dim(V_{111}) - \delta$, and plugging these into Equation (6.2) gives us $\Lambda_{\text{even}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i\in\mathcal{U}}) \geq 0$ for some 2^7-dimensional vector Λ_{even}, as desired; after applying these steps one obtains Equation (6.8) below. When $\{V_i\}_{i\in\mathcal{U}}$ are one-dimensional subspaces constituting a representation of the non-Fano matroid over a field of characteristic $\neq 2$, it is easy to check that all of the δ terms appearing in (6.7) are zero. So, the inequality states that $\dim(V_{110}, V_{101}, V_{011}) \leq 2\dim(V_{111})$, whereas we know that $\dim(V_{110}, V_{101}, V_{011}) = 3\dim(V_{111})$ for the non-Fano matroid. Consequently $\Lambda_{\text{odd}} \cdot \vec{\mathbf{r}}(\mathcal{F}) < 0$.

For completeness, the inequality $\Lambda_{\text{even}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i\in\mathcal{U}}) \geq 0$ is written explicitly as follows.

$$\begin{aligned}&2\dim(V_{100}) + 2\dim(V_{010}) + 3\dim(V_{001}) + 11\dim(V_{111}) \\ &+3\dim(V_{100}, V_{010}) + 2\dim(V_{100}, V_{001}) + 2\dim(V_{010}, V_{001}) \\ &-\dim(V_{100}, V_{111}) - \dim(V_{010}, V_{111}) - \dim(V_{001}, V_{111}) - 4\dim(V_{100}, V_{010}, V_{001}) \\ &-3\dim(V_{111}, V_{100}, V_{010}) - 3\dim(V_{111}, V_{100}, V_{001}) - 3\dim(V_{111}, V_{010}, V_{001}) \\ &+\dim(V_{110}, V_{100}, V_{010}) + \dim(V_{101}, V_{100}, V_{001}) + \dim(V_{011}, V_{010}, V_{001}) \\ &+\dim(V_{110}, V_{111}, V_{001}) + \dim(V_{101}, V_{111}, V_{010}) + \dim(V_{011}, V_{111}, V_{100}) \\ &-\dim(V_{110}, V_{101}, V_{011}) + \dim(V_{111}, V_{100}, V_{010}, V_{001}) \geq 0\,.\end{aligned} \tag{6.8}$$

This concludes the proof of the theorem. ■

6.2 Linear codes over fields of odd characteristic

The following lemma and theorem, which are analogues of Lemma 6.3 and Theorem 6.4 from Section 6.1, provide an inequality for fields with odd characteristic.

Lemma 6.5 (Conditional Odd Characteristic Inequality)**.** *Suppose* $\{V_i\}_{i\in\mathcal{U}}$ *are 7 subspaces of a vector space over* $\mathbb{F}$ *such that* $\text{char}(\mathbb{F}) \neq 2$ *and*

(i) $\dim(\{V_i\}_{i\in\mathcal{O}}) = \dim(\{V_i\}_{i\in\mathcal{B}})$

(ii) $\dim(V_i, V_j, V_k) = \dim(V_i) + \dim(V_j) + \dim(V_k)\ \forall i,j,k \in \mathcal{O}$

(iii) $\dim(V_i, V_j, V_{i+j}) = \dim(V_i, V_j)\ \forall i,j \in \mathcal{B}$

(iv) $\dim(V_i, V_j, V_{111}) = \dim(V_i, V_j)\ \forall i,j : i+j = 111$

Then $\dim(V_{110}, V_{101}, V_{011}) \geq 3\dim(V_{111})$.

Proof. Just as in the proof of Lemma 6.3 we apply the result of Lemma 6.2, but now with $\mathrm{char}(\mathbb{F}) \neq 2$. Hypotheses (i) and (iii) imply that all 7 subspaces are contained in the span of $V_{100}, V_{010}, V_{001}$, and hypothesis (ii) implies that those three subspaces are mutually complementary, and that V_{111} is complementary to the sum of any two of them. Thus, Lemma 6.2 implies that $\dim(V_{110}, V_{101}, V_{011}) = 3\dim(W)$. Now we aim to show that hypotheses (iii) and (iv) imply that V_{110} contains $\pi_{12}(V_{111})$, and similarly for V_{101}, V_{011}. This will imply that $\dim(V_{110}, V_{101}, V_{011}) \geq \dim(\pi_{12}(W), \pi_{23}(W), \pi_{13}(W)) = 3\dim(W)$ as desired.

It remains for us to justify the claim that V_{110} contains $\pi_{12}(V_{111})$. Suppose (x,y,z) belongs to V_{111}, where we use (x,y,z) as an alternate notation for $x+y+z$ such that x belongs to V_{100}, y belongs to V_{010}, z belongs to V_{001}. We know from hypothesis (iv) that V_{111} is contained in $V_{001} + V_{110}$. So write $x+y+z = a+b$ where a is in V_{001} and b is in V_{110}. We know from hypothesis (iii) that V_{110} is contained in $V_{100}+V_{010}$, so write $b = c+d$ where c is in V_{100} and d is in V_{010}. Then $x+y+z = c+d+a$, and both sides are a sum of three vectors, the first belonging to V_{100}, the second to V_{010}, the third to V_{001}. Since those three vector spaces are mutually complementary, the representation of another vector as a sum of vectors from each of them is unique. So $x = c, y = d, z = a$. This means that $x + y = c + d = \pi_{12}(x,y,z)$. Recall that $c + d$ is in V_{110}. As (x,y,z) was an arbitrary element of V_{111}, we have shown that V_{110} is contained in $\pi_{12}(V_{111})$. ■

Theorem 6.6 (Odd Characteristic Inequality)**.** *There exists a* 2^7*-dimensional vector* Λ_{odd} *such that for any 7 subspaces* $\{V_i\}_{i\in\mathcal{U}}$ *of a vector space over* $\mathbb{F}$ *with* $\mathrm{char}(\mathbb{F}) \neq 2$*,*

$$\Lambda_{\mathrm{odd}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i\in\mathcal{U}}) \geq 0 \text{ and } \Lambda_{\mathrm{odd}} \cdot \vec{\mathbf{r}}(\mathcal{F}) < 0\,.$$

Proof. Let $\{V_i\}_{i\in\mathcal{U}}$ be seven subspaces of a vector space V over $\mathbb{F}$ such that $\mathrm{char}(\mathbb{F}) \neq 2$. Just as in the proof Theorem 6.4, we will modify the subspaces $\{V_i\}_{i\in\mathcal{U}}$ into $\{V_i'\}_{i\in\mathcal{U}}$ that satisfy the conditions of Lemma 6.5, starting with $\{V_i'\}_{i\in\mathcal{U}} = \{V_i\}_{i\in\mathcal{U}}$. We again delete vectors of a certain type in an iterative fashion. The order of the deletions within each step is arbitrary.

Step 1: Vectors in V_{111}' but not in $\sum_{i\in\mathcal{B}} V_i'$ from V_{111}'.

Step 2: (a) Vectors in $V_{100}' \cap V_{010}'$ from V_{010}'.
(b) Vectors in $V_{001}' \cap (V_{100}' + V_{010}')$ from V_{001}'.
(c) Vectors in $V_{111}' \cap (V_{100}' + V_{010}')$ from V_{111}'.
(d) Vectors in $V_{111}' \cap (V_{010}' + V_{001}')$ from V_{111}'.
(e) Vectors in $V_{111}' \cap (V_{100}' + V_{001}')$ from V_{111}'.

Step 3: Vectors in V_{i+j}' but not in $V_i' + V_j'$ for $i,j \in \mathcal{B}$ from V_{i+j}'.

Step 4: Vectors in V'_{111} but not in $V'_i + V'_j$ for $i, j : i + j = 111$ from V'_{111}.

The first two steps in this sequence of deletions, along with the first two conditions in Lemma 6.5 are identical to those in the even characteristic case. Thus, by arguments from the proof Theorem 6.4 we have that by the end of step (2) conditions (i), (ii) are satisfied. Step (3) is almost identical to the same step in the even characteristic case; the difference is that now we only perform the step for pairs $i, j \in \mathcal{B}$ rather than all pairs $i, j \in \mathcal{O}$. As before, at the end of step (3) condition (iii) is satisfied, and since the step does not modify V'_i for any $i \in \mathcal{O}$, it does not cause either of conditions (i), (ii) to become violated. Step (4) ensures condition (iv), so it remains to show that step (4) preserves conditions (i)–(iii). Step (4) only modifies V'_{111} so it doesn't change $\sum_{i\in\mathcal{B}} V'_i$, therefore preserving (i). It preserves (ii) because if three subspaces are mutually complementary, they remain mutually complementary after deleting a vector from one of them. It preserves (iii) because (iii) does not involve V'_{111}, which is the only subspace that changes during step (4).

Now, by Lemma 6.3 we have that

$$3\dim(V'_{111}) \leq \dim(V'_{110}, V'_{101}, V'_{011}). \tag{6.9}$$

As in the proof of Theorem 6.4, let

$$\begin{aligned}
\delta &= \dim(V_{111}, \{V_i\}_{i\in\mathcal{B}}) - \dim(\{V_i\}_{i\in\mathcal{B}}) \\
\delta[i|j,k] &= \dim(V_i, V_j, V_k) - \dim(V_j, V_k) \\
\delta[i;j] &= \dim(V_i \cap V_j) = \dim(V_i) + \dim(V_j) - \dim(V_i, V_j) \\
\delta[i;j,k] &= \dim(V_i \cap (V_j + V_k)) = \dim(V_i) + \dim(V_j, V_k) - \dim(V_i, V_j, V_k)
\end{aligned}$$

Observe that we only reduce the size of subspaces, so $\dim(V'_{110}, V'_{101}, V'_{011}) \leq \dim(V_{110}, V_{101}, V_{011})$.

It remains to get a lower bound on $\dim(V'_{111})$ in terms of dimensions of subsets of $\{V_i\}_{i\in\mathcal{U}}$. We do this by giving an upper bound on the number of vectors we delete from V'_{111} in terms of the δ terms we defined above. Step (1) deletes δ vectors. Steps (2a) and (2b) delete nothing from V'_{111}, and at the end of (2a)–(2b) we have

$$\dim(V'_{100}) = \dim(V_{100}) \tag{6.10}$$

$$\dim(V'_{010}) = \dim(V_{010}) - \delta[100;010] \tag{6.11}$$

$$\dim(V'_{001}) = \dim(V_{001}) - \delta[001;100,010] \tag{6.12}$$

Steps (2c)–(2e) delete at most $\sum_{i,j\in\mathcal{B}} \delta[111;i,j]$ vectors from V'_{111}, and they do not change any of the other subspaces.

In step (3) no vectors are deleted from V'_{111}, but we will still need an upper bound on the number of vectors deleted in this step since it will influence our upper bound on the number of vectors deleted from V'_{111} in step (4). If no deletions took place prior to step (3), then for all $i, j \in \mathcal{B}$ exactly $\delta[i+j|i,j]$ vectors would be deleted from V'_{i+j} during step (3). However, if $\dim(V'_i, V'_j) < \dim(V_i, V_j)$, then we must adjust our estimate of the number of deleted vectors to account for this difference. Steps (1) and (2a) cannot change $\dim(V'_i, V'_j)$ for any $i, j \in \mathcal{B}$, but step (2b) reduces each of $\dim(V'_{001}, V'_{100})$ and $\dim(V'_{001}, V'_{010})$ by at most $\delta[001;100,010]$. Therefore, at

the end of step (3) we have

$$\dim(V'_{110}) = \dim(V_{110}) - \delta[110|100,010] \tag{6.13}$$
$$\dim(V'_{101}) \geq \dim(V_{101}) - \delta[101|100,001] - \delta[001;100,010] \tag{6.14}$$
$$\dim(V'_{011}) \geq \dim(V_{011}) - \delta[011|010,001] - \delta[001;100,010] \tag{6.15}$$

If no deletions took place prior to step (4), then the number of vectors we would need to delete from V'_{111}, to make it a subspace of $V'_i + V'_j$, would be at most $\delta[111|i,j]$. As before, we need to adjust this bound to account for the potential difference in dimension between $V_i + V_j$ and $V'_i + V'_j$. Using the upper bound

$$\dim(V_i + V_j) - \dim(V'_i + V'_j) \leq \big[\dim(V_i) - \dim(V'_i)\big] + \big[\dim(V_j) - \dim(V'_j)\big]\,,$$

which is valid for any four subspaces V_i, V_j, V'_i, V'_j satisfying $V'_i \subseteq V_i$, $V'_j \subseteq V_j$, we find that the number of extra vectors deleted from V'_{111} in step (4) because of differences in dimension between $V'_i + V'_j$ and $V_i + V_j$ (for some $i, j \in \mathcal{U}$, $i + j = 111$), is at most

$$\sum_{i \in \mathcal{U} \setminus \{111\}} \dim(V_i) - \dim(V'_i) \leq \delta[100;010] + 3\delta[001;100,010] + \sum_{i,j \in \mathcal{B}} \delta[i+j|i,j],$$

where the first inequality follows by combining equations (6.10)–(6.15).

We now sum up our upper bounds on the number of vectors deleted from V'_{111} in steps (1)–(4) combined, to find that

$$\dim(V'_{111}) \geq \dim(V_{111}) - \delta - \sum_{i,j \in \mathcal{B}} \delta[111;i,j] - \delta[100;010] - 3\delta[001;100,010] - \sum_{i,j \in \mathcal{B}} \delta[i+j|i,j]. \tag{6.16}$$

Expanding out all of the δ terms, combining with the upper bound on $\dim(V'_{111})$, and plugging these into Equation (6.9) gives us $\Lambda_{\text{odd}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i \in \mathcal{U}}) \geq 0$ for some 2^7-dimensional vector Λ_{odd}, as desired; after applying these steps one obtains Equation (6.17) below. When $\{V_i\}_{i \in \mathcal{U}}$ are one-dimensional subspaces constituting a representation of the Fano matroid over a field of characteristic 2, it is easy to check that all of the δ terms appearing in (6.16) are zero. So, the inequality states that $\dim(V_{110}, V_{101}, V_{011}) \geq 3\dim(V_{111})$, whereas we know that $\dim(V_{110}, V_{101}, V_{011}) = 2\dim(V_{111})$ for the Fano matroid. Consequently $\Lambda_{\text{odd}} \cdot \vec{\mathbf{r}}(\mathcal{F}) < 0$.

For completeness, the inequality $\Lambda_{\text{odd}} \cdot \vec{\mathbf{d}}(\{V_i\}_{i \in \mathcal{U}}) \geq 0$ is written explicitly as follows.

$$\begin{aligned}
&3\dim(V_{100}) + 3\dim(V_{010}) + 9\dim(V_{001}) + 6\dim(V_{111}) + 6\dim(V_{100}, V_{010}) - 12\dim(V_{100}, V_{010}, V_{001}) \\
&+3\dim(V_{110}, V_{100}, V_{010}) + 3\dim(V_{101}, V_{100}, V_{001}) + 3\dim(V_{011}, V_{010}, V_{001}) \\
&-3\dim(V_{111}, V_{100}, V_{010}) - 3\dim(V_{111}, V_{100}, V_{001}) - 3\dim(V_{111}, V_{010}, V_{001}) \\
&+3\dim(V_{111}, V_{100}, V_{010}, V_{001}) + \dim(V_{110}, V_{101}, V_{011}) \geq 0\,.
\end{aligned} \tag{6.17}$$

This completes the proof of the theorem. ■

6.3 Polynomial separation between the linear and non-linear rates

The following pair of lemmas shows how to take a single linear constraint, such as one of those whose existence is asserted by Theorems 6.4 and 6.6, and transform it into a tight homomorphic constraint schema. To state the lemmas, we must first define the set of vectors $D_{\mathbb{F}}(K) \subset \mathbb{R}^{\mathcal{P}(K)}$, for any index set K and field $\mathbb{F}$, to be the set of all vectors $\vec{\mathbf{d}}(\{V_k\}_{k \in K})$, where $\{V_k\}_{k \in K}$ runs through all K-indexed tuples of finite-dimensional vector spaces over $\mathbb{F}$.

Lemma 6.7 (Tightening Modification)**.** *Suppose I is any index set, e is an element not in I, and $J = I \cup \{e\}$. There exists an explicit linear transformation from $\mathbb{R}^{\mathcal{P}(J)}$ to $\mathbb{R}^{\mathcal{P}(I)}$, represented by a matrix B, such that:*

(i) $B \cdot D_{\mathbb{F}}(J) \subseteq D_{\mathbb{F}}(I)$ *for every field* $\mathbb{F}$.

(ii) $B\mathbf{1} = B\mathbf{1}_j = 0$ *for all* $j \in J$.

(iii) If M is a matroid with ground set I and the intersection of all matroid bases of M is the empty set, then $B\vec{\mathbf{r}}(M + e) = \vec{\mathbf{r}}(M)$, where $M + e$ denotes the matroid obtained by adjoining a rank-zero element to M.

Proof. If U is any vector space with a J-tuple of subspaces $\{U_j\}_{j\in J}$, then there is a quotient map π from U to $V = U/U_e$, and we can form an I-tuple of subspaces $\{V_i\}_{i\in I}$ by specifying that $V_i = \pi(U_i)$ for all $i \in I$. The dimension vectors $\vec{\mathbf{u}} = \vec{\mathbf{d}}(\{U_j\})$ and $\vec{\mathbf{v}} = \vec{\mathbf{d}}(\{V_i\})$ are related by an explicit linear transformation. In fact, for any subset $S \subseteq I$, if we let U_S, V_S denote the subspaces of U, V spanned by $\{U_i\}_{i\in S}$ and $\{V_i\}_{i\in S}$, respectively, then π maps $U_S + U_e$ onto V_S with kernel U_e, and this justifies the formula

$$\mathbf{v}_S = \mathbf{u}_{S\cup\{e\}} - \mathbf{u}_{\{e\}}.$$

Thus, $\mathbf{v} = B_0\mathbf{u}$, where B_0 is the matrix

$$(B_0)_{ST} = \begin{cases} 1 & \text{if } T = S \cup \{e\} \\ -1 & \text{if } T = \{e\} \\ 0 & \text{otherwise,} \end{cases} \tag{6.18}$$

and therefore $B_0 \cdot D_{\mathbb{F}}(J) \subseteq D_{\mathbb{F}}(I)$.

Similarly, if U is any vector space with an I-tuple of subspaces $\{U_i\}_{i\in I}$ and k is any element of I, we can define $U_{-k} \subseteq U$ to be the linear subspace spanned by $\{U_i\}_{i\neq k}$, and we can let $\pi : U \to U_{-k}$ be any linear transformation whose restriction to U_{-k} is the identity. The restriction of π to U_k has kernel W_k of dimension $\dim(W_k) = \dim(\{U_i\}_{i\in I}) - \dim(\{U_i\}_{i\in I, i\neq k})$. As before, let $V_i = \pi(U_i)$ for all $i \in I$, let U_S, V_S denote the subspaces of U, V spanned by $\{U_i\}_{i\in S}$ and $\{V_i\}_{i\in S}$, and let $\vec{\mathbf{u}} = \vec{\mathbf{d}}(\{U_i\}), \vec{\mathbf{v}} = \vec{\mathbf{d}}(\{V_i\})$. If $k \notin S$ then $V_S = U_S$ and $\mathbf{v}_S = \mathbf{u}_S$, while if $k \in S$ then U_S contains W_k, the linear transformation π maps U_S onto V_S with kernel W_k, and $\mathbf{v}_S = \mathbf{u}_S - \dim(W_k) = \mathbf{u}_S - \mathbf{u}_I + \mathbf{u}_{I\setminus\{k\}}$. Thus, $\mathbf{v} = B_k\mathbf{u}$, where B_k is the matrix

$$(B_k)_{ST} = \begin{cases} 1 & \text{if } T = S \\ 1 & \text{if } k \in S \text{ and } T = I \setminus \{k\} \\ -1 & \text{if } k \in S \text{ and } T = I \\ 0 & \text{otherwise.} \end{cases} \tag{6.19}$$

and therefore $B_k \cdot D_{\mathbb{F}}(I) \subseteq D_{\mathbb{F}}(I)$.

Now assume without loss of generality that $I = \{1, 2, \ldots, n\}$ and let $B = B_nB_{n-1}\cdots B_1B_0$. We have seen that $B \cdot D_{\mathbb{F}}(J) \subseteq D_{\mathbb{F}}(I)$. From (6.18) one can see that $B_0\mathbf{1} = B_0\mathbf{1}_e = 0$ and that for every $k \in I$, $B_0\mathbf{1}_k = \mathbf{1}_k$. (Here, it is important to note that $\mathbf{1}_k$ on the left side refers to a vector in $\mathbb{R}^{\mathcal{P}(J)}$ and on the right side it refers to a vector in $\mathbb{R}^{\mathcal{P}(I)}$.) Furthermore, from (6.19) one can see that $B_k\mathbf{1}_k = 0$ and that $B_k\mathbf{1}_i = \mathbf{1}_i$ for all $i \neq k$. Thus, when we left-multiply a vector $\vec{\mathbf{w}} \in \{\mathbf{1}\} \cup \{\mathbf{1}_j\}_{j\in J}$ by the matrix B, one of the following things happens. If $\vec{\mathbf{w}}$ is equal to $\mathbf{1}$ or $\mathbf{1}_e$

then $B_0\vec{\mathbf{w}} = 0$ hence $B\vec{\mathbf{w}} = 0$. Otherwise, $\vec{\mathbf{w}} = \mathbf{1}_k \in \mathbb{R}^{\mathcal{P}(J)}$ for some $k \in I$, $B_0\vec{\mathbf{w}} = \mathbf{1}_k \in \mathbb{R}^{\mathcal{P}(I)}$, and as we proceed to left-multiply $\mathbf{1}_k$ by $B_1, B_2, \ldots$, it is fixed by B_i $(i < k)$ and annihilated by B_k, so once again $B\vec{\mathbf{w}} = 0$. This confirms assertion (ii) of the lemma.

Finally, if $M, M + e$ are matroids satisfying the hypotheses of assertion (iii), then for every set $S \subseteq I$ we have $r(S \cup \{e\}) - r(\{e\}) = r(S)$ and hence $B_0\vec{\mathbf{r}}(M+e) = \vec{\mathbf{r}}(M)$. For any $k \in I$ our assumption on M implies that it has a matroid basis disjoint from $\{k\}$, and hence that $r(I \setminus \{k\}) = r(I)$. Inspecting (6.19), we see that this implies $B_k\vec{\mathbf{r}}(M) = \vec{\mathbf{r}}(M)$ for all $k \in I$, and hence $B\vec{\mathbf{r}}(M+e) = \vec{\mathbf{r}}(M)$ as desired. ∎

Lemma 6.8 (Homomorphic Schema Extension)**.** *Let I be an index set, and let $\vec{\alpha} \in \mathbb{R}^{\mathcal{P}(I)}$ be a vector such that $\vec{\alpha}^{\mathsf{T}}\vec{\mathbf{d}} \geq 0$ for all $\vec{\mathbf{d}} \in D_{\mathbb{F}}(I)$. Then there is a homomorphic constraint schema (Q, A) such that $\vec{\alpha}^{\mathsf{T}}$ is a row of the matrix $A(I)$, and for every index set K and vector $\vec{\mathbf{d}} \in D_{\mathbb{F}}(K)$, $A(K)\vec{\mathbf{d}} \geq 0$. If $\vec{\alpha}^{\mathsf{T}}\mathbf{1} = \vec{\alpha}^{\mathsf{T}}\mathbf{1}_i = 0$ for all $i \in I$, then the constraint schema (Q, A) is tight.*

Proof. For any index set J, let $Q(J)$ be the set of all Boolean lattice homomorphisms from $\mathcal{P}(I)$ to $\mathcal{P}(J)$. If $h : \mathcal{P}(J) \to \mathcal{P}(K)$ is another Boolean lattice homomorphism, the mapping h_* is defined by function composition, i.e. $h_*(q) = h \circ q$.

To define the constraint matrix $A(J)$ associated to an index set J, we do the following. A row of $A(J)$ is indexed by a Boolean lattice homomorphism $q : \mathcal{P}(I) \to \mathcal{P}(J)$ and we define the entries of that row by

$$A(J)_{qS} = \sum_{\substack{T \in \mathcal{P}(I) \\ q(T) = S}} \alpha_T. \tag{6.20}$$

This defines a homomorphic constraint schema, because if $h : \mathcal{P}(J) \to \mathcal{P}(K)$ is any Boolean lattice homomorphism and $R = A(K)^{\mathsf{T}}Q_h$, $R' = P_h A(J)^{\mathsf{T}}$, then recalling the definitions of P_h, Q_h we find that

$$R_{Sq} = (A(K)^{\mathsf{T}})_{S,h_*(q)} = A(K)_{h\circ q,S} = \sum_{\substack{T \in \mathcal{P}(I) \\ h(q(T)) = S}} \alpha_T$$

$$R'_{Sq} = \sum_{S' : h(S') = S} (A(J)^{\mathsf{T}})_{S',q} = \sum_{S' : h(S') = S} \sum_{\substack{T \in \mathcal{P}(I) \\ q(T) = S'}} \alpha_T$$

and the right-hand sides of the two lines are clearly equal.

To prove that $A(K)\vec{\mathbf{d}} \geq 0$ for every $\vec{\mathbf{d}} \in D_{\mathbb{F}}(K)$, we reason as follows. It suffices to take a single row of the constraint matrix, indexed by homomorphism $q : I \to K$, and to prove that

$$\sum_{S \in \mathcal{P}(K)} A(K)_{qS}\vec{\mathbf{d}}_S \geq 0.$$

Using the definition of the constraint matrix entries, this can be rewritten as

$$\sum_{T \in \mathcal{P}(I)} \alpha_T \vec{\mathbf{d}}_{q(T)} \geq 0. \tag{6.21}$$

Let $\{V_k\}_{k\in K}$ be a K-tuple of vector spaces such that $\vec{\mathbf{d}} = \vec{\mathbf{d}}(\{V_k\}_{k\in K})$. Define an I-tuple of vector spaces $\{U_i\}_{i\in I}$ by setting U_i to be the span of $\{V_k\}_{k\in q(\{i\})}$. By our hypothesis on $\vec{\alpha}$,

$$\sum_{T \in \mathcal{P}(I)} \alpha_T \dim(\{U_i\}_{i\in T}) \geq 0.$$

The left side is equal to the left side of (6.21).

Finally, suppose that $\vec{\alpha}^{\mathsf{T}}\mathbf{1} = \vec{\alpha}^{\mathsf{T}}\mathbf{1}_i$ for all $i \in I$. For any index set J, we prove that $A(J)\mathbf{1} = 0$ by calculating the component of $A(J)\mathbf{1}$ indexed by an arbitrary Boolean lattice homomorphism $q : \mathcal{P}(I) \to \mathcal{P}(J)$.

$$\sum_{S \in \mathcal{P}(J)} A(J)_{qS} = \sum_{S \in \mathcal{P}(J)} \sum_{\substack{T \in \mathcal{P}(I) \\ q(T) = S}} \alpha_T = \sum_{T \in \mathcal{P}(I)} \alpha_T = \vec{\alpha}^{\mathsf{T}}\mathbf{1} = 0$$

The proof that $A(J)\mathbf{1}_j = 0$ for all $j \in J$ similarly calculates the component of $A(J)\mathbf{1}_j$ indexed by an arbitrary q.

$$\sum_{\substack{S \in \mathcal{P}(J) \\ j \in S}} A(J)_{qS} = \sum_{\substack{S \in \mathcal{P}(J) \\ j \in S}} \sum_{\substack{T \in \mathcal{P}(I) \\ q(T) = S}} \alpha_T = \sum_{\substack{T \in \mathcal{P}(I) \\ j \in q(T)}} \alpha_T$$

At this point the argument splits into three cases. If $j \in q(\emptyset)$ then the right side is $\vec{\alpha}^{\mathsf{T}}\mathbf{1}$, which equals 0. If $j \notin q(J)$ then the right side is an empty sum and clearly equals 0. If $j \notin q(\emptyset)$ but $j \in q(J)$, then there is a unique $i \in I$ such that $j \in q(\{i\})$. Indeed, if j belongs to $q(\{i\})$ and $q(\{i'\})$, then j belongs to $q(\{i\}) \cap q(\{i'\}) = q(\{i\} \cap \{i'\})$, implying that $\{i\} \cap \{i'\}$ is non-empty and that $i = i'$. The right side of the equation above is thus equal to $\vec{\alpha}^{\mathsf{T}}\mathbf{1}_i$, which equals 0. ∎

Finally, before proving Theorem 6.1, it will be useful to describe the following simple operation for combining constraint schemas.

Definition 6.9. The *disjoint union* of two constraint schemas (Q_1, A_1) and (Q_2, A_2) is the constraint schema which associates to every index set I the disjoint union $\mathcal{Q}(I) = \mathcal{Q}_1(I) \sqcup \mathcal{Q}_2(I)$ and the constraint matrix $A(I)$ given by

$$A(I)_{qS} = \begin{cases} A_1(I)_{qS} & \text{if } q \in \mathcal{Q}_1(I) \\ A_2(I)_{qS} & \text{if } q \in \mathcal{Q}_2(I). \end{cases}$$

For a homomorphism $h : \mathcal{P}(I) \to \mathcal{P}(J)$, the function $h_* : \mathcal{Q}_1(I) \sqcup \mathcal{Q}_2(I) \to \mathcal{Q}_1(J) \sqcup \mathcal{Q}_2(J)$ is defined by combining $\mathcal{Q}_1(I) \xrightarrow{h_*} \mathcal{Q}_1(J)$ and $\mathcal{Q}_2(I) \xrightarrow{h_*} \mathcal{Q}_2(J)$ in the obvious way.

Lemma 6.10. *The disjoint union of two tight constraint schemas is tight, and the disjoint union of two homomorphic constraint schemas is homomorphic.*

Proof. For all index sets I and vectors $v \in \mathbb{R}^{\mathcal{P}(I)}$, the constraint matrix of the disjoint union satisfies

$$A(I)\mathbf{v} = \begin{pmatrix} A_1(I) \\ A_2(I) \end{pmatrix} \mathbf{v} = \begin{pmatrix} A_1(I)\mathbf{v} \\ A_2(I)\mathbf{v} \end{pmatrix}$$

so if both constraint schemas are tight then so is their disjoint union. If $h : \mathcal{P}(I) \to \mathcal{P}(J)$ is a Boolean lattice homomorphism then let Q_{1h}, Q_{2h}, Q_h denote the matrices representing the induced linear transformations $\mathbb{R}^{\mathcal{Q}_1(I)} \to \mathbb{R}^{\mathcal{Q}_1(J)}$, $\mathbb{R}^{\mathcal{Q}_2(I)} \to \mathbb{R}^{\mathcal{Q}_2(J)}$, and $\mathbb{R}^{\mathcal{Q}(I)} \to \mathbb{R}^{\mathcal{Q}(J)}$, respectively. If both constraint schemas are homomorphic, then

$$\begin{aligned} A(J)^{\mathsf{T}}Q_h = \begin{pmatrix} A_1(J)^{\mathsf{T}} & A_2(J)^{\mathsf{T}} \end{pmatrix} \begin{pmatrix} Q_{1h} & 0 \\ 0 & Q_{2h} \end{pmatrix} &= \begin{pmatrix} A_1(J)^{\mathsf{T}}Q_{1h} & A_2(J)^{\mathsf{T}}Q_{2h} \end{pmatrix} \\ &= \begin{pmatrix} P_h A_1(J)^{\mathsf{T}} & P_h A_2(J)^{\mathsf{T}} \end{pmatrix} = P_h A(J)^{\mathsf{T}}, \end{aligned}$$

which confirms that the disjoint union is homomorphic. ∎

The proof of Theorem 6.1 now follows by combining earlier results.

Proof of Theorem 6.1 (Separation Theorem). The fact that $\beta(G_{\mathcal{F}}) = \beta(G_{\mathcal{N}}) = 4$ is an immediate consequence of Theorem 5.4. The submultiplicativity of β under the lexicographic product (Theorem 2.2) then implies that $G = G_{\mathcal{F}} \bullet G_{\mathcal{N}}$ satisfies $\beta(G^{\bullet n}) \leq (4 \cdot 4)^n = 16^n$. A lower bound of the form $\beta(G^{\bullet n}) \geq 16^n$ is a consequence of Proposition 5.2 which implied that $b(G_{\mathcal{F}}) = b(G_{\mathcal{N}}) = 4$, from which it follows by the supermultiplicativity of b under lexicographic products (Theorem 3.4) that $16^n \leq b(G^{\bullet n}) \leq \beta(G^{\bullet n})$. Combining these upper and lower bounds, we find that $\beta(G^{\bullet n}) = 16^n$.

It is worth noting, incidentally, that although each of $G_{\mathcal{F}}$, $G_{\mathcal{N}}$ individually has a linear solution over the appropriate field, the index code for $G = G_{\mathcal{F}} \bullet G_{\mathcal{N}}$ implied by the proof of Theorem 2.2 — which concatenates these two linear codes together by composing them with an arbitrary one-to-one mapping from a mod-2 vector space to a mod-p vector space (p odd) — is highly nonlinear, and not merely a side-by-side application of two linear codes.

To establish the lower bound on $\lambda^{\mathbb{F}}(G^{\bullet n})$, we distinguish two cases, $\mathrm{char}(\mathbb{F}) = 2$ and $\mathrm{char}(\mathbb{F}) \neq 2$, and in both cases we prove $\lambda^{\mathbb{F}}(G^{\bullet n}) \geq (16+\varepsilon)^n$ using the LP in Figure 1 with a tight homomorphic constraint schema supplying the constraint matrix A. We use different constraint schemas in the two cases but the constructions are nearly identical. Let M denote the matroid $\mathcal{N}$ if $\mathrm{char}(\mathbb{F}) = 2$, and let $M = \mathcal{F}$ if $\mathrm{char}(\mathbb{F}) \neq 2$. In both cases, we will let $M + e$ denote the matroid obtained by adjoining a rank-zero element to M, and we will denote the ground sets of M, $M + e$ by I, J, respectively. Recall the vectors $\Lambda_{\mathrm{even}}, \Lambda_{\mathrm{odd}} \in \mathbb{R}^{\mathcal{P}(I)}$ from Theorems 6.4 and 6.6. Let $\Lambda = \Lambda_{\mathrm{even}}$ if $\mathrm{char}(\mathbb{F}) = 2$, $\Lambda = \Lambda_{\mathrm{odd}}$ if $\mathrm{char}(\mathbb{F}) \neq 2$. By Theorems 6.4 and 6.6, $\Lambda \cdot \vec{\mathbf{r}}(M) < 0$, a fact that we will be using later.

Recall the linear transformation $B : \mathbb{R}^{\mathcal{P}(J)} \to \mathbb{R}^{\mathcal{P}(I)}$ from Lemma 6.7, and let

$$\vec{\alpha} = B^{\mathsf{T}}\Lambda.$$

For any $\vec{\mathbf{d}} \in D_{\mathbb{F}}(J)$ we have $\vec{\alpha}^{\mathsf{T}}\vec{\mathbf{d}} = \Lambda^{\mathsf{T}}B\vec{\mathbf{d}} \geq 0$, since $B\vec{\mathbf{d}} \in D_{\mathbb{F}}(I)$ and $\Lambda \cdot \vec{\mathbf{v}} \geq 0$ for all $\vec{\mathbf{v}} \in D_{\mathbb{F}}(I)$. The equations $B\mathbf{1} = B\mathbf{1}_j = 0$ for all $j \in J$ imply that $\vec{\alpha}^{\mathsf{T}}\mathbf{1} = \vec{\alpha}^{\mathsf{T}}\mathbf{1}_j = 0$. Applying Lemma 6.8 to obtain a tight homomorphic constraint schema from $\vec{\alpha}$, and taking its disjoint union with the submodularity constraint schema, we arrive at a tight homomorphic constraint schema (Q, A) such that every vector $\vec{\mathbf{d}} \in D_{\mathbb{F}}(K)$, for every index set K, satisfies the system of inequalities $A(K)\vec{\mathbf{d}} \geq 0$.

Consider the LP in Figure 1, instantiated with constraint schema (Q, A). We claim that its optimal solution $b^{\mathbb{F}}(G)$, for any index coding problem G, satisfies $b^{\mathbb{F}}(G) \leq \lambda^{\mathbb{F}}(G)$. To prove this we proceed as in the proof of Theorem 3.1: consider any linear index code over $\mathbb{F}$ with message alphabet Σ, sample each message independently and uniformly at random, and consider the input messages and the broadcast message as random variables, with $H(S)$ denoting the joint entropy of a subset S of these random variables. The set of random variables is indexed by $K = V(G) \cup \{e\}$, where $V(G)$ denotes the set of messages in G and e is an extra element of K corresponding to the broadcast message. Letting M_k, for $k \in K$, denote the matrix representing the linear transformation defined by the k^{th} random variable, and letting U_k denote the row space of M_k, we have $H(S) = \log|\mathbb{F}| \cdot \dim(\{U_k\}_{k \in S})$ for every $S \subseteq K$. For $S \subseteq V(G)$ let

$$z_S = \frac{H(S \cup \{e\})}{\log|\Sigma|} = \left(\frac{\log|\mathbb{F}|}{\log|\Sigma|}\right) \dim(U_e,\, \{U_i\}_{i \in S}).$$

We aim to show that z satisfies the constraints of the LP, implying that $b^{\mathbb{F}}(G) \leq z_{\emptyset} = \frac{\log|\Sigma_P|}{\log|\Sigma|}$ and consequently (since the linear index code over $\mathbb{F}$ was arbitrary) that $b^{\mathbb{F}}(G) \leq \lambda^{\mathbb{F}}(G)$. The proof

of Theorem 3.1 already established that $z_{V(G)} = |V(G)|$ and that $z_T - z_S \le c_{ST}$ for all $S \subset T$, so we need only show that $Az \ge 0$. As in the proof of Lemma 6.7, we let π denote the quotient map from $U = \Sigma^{V(G)}$ to U/U_e, we define $V_i = \pi(U_i)$, and we observe that for all $S \subseteq V(G)$, $\dim(\{V_i\}_{i\in S}) = \dim(U_e,\, \{U_i\}_{i\in S}) - \dim(U_e)$. This implies that

$$z - z_\emptyset \mathbf{1} = \left(\frac{\log|\mathbb{F}|}{\log|\Sigma|}\right)\vec{\mathbf{d}},$$

where $\vec{\mathbf{d}} = \vec{\mathbf{d}}(\{V_i\})$. Our construction of A implies that $A\vec{\mathbf{d}} \ge 0$ and that $A\mathbf{1} = 0$, hence $Az \ge 0$.

It remains to show that $b^{\mathbb{F}}(G) > 16$, from which the claimed lower bound follows by supermultiplicativity. Since our constraint schema includes submodularity, we have $b^{\mathbb{F}}(G_{\mathcal{F}}) \ge b(G_{\mathcal{F}}) = 4$ and $b^{\mathbb{F}}(G_{\mathcal{N}}) \ge b(G_{\mathcal{N}}) = 4$, so we need only show that one of these inequalities is strict, and we accomplish this using Theorem 5.5. Specifically, we show that the matrix $A = A(I)$ has a row indexed by some $q \in \mathcal{Q}(I)$, such that $(A\,\vec{\mathbf{r}}(M))_q < 0$. Recall that $\mathcal{Q}(I)$ is the set of Boolean lattice homomorphisms from $\mathcal{P}(J)$ to $\mathcal{P}(I)$, where $J = I \cup \{e\}$. Let q be the homomorphism that maps $\{e\}$ to $\emptyset$ and $\{i\}$ to itself for every $i \in I$. To prove that $(A\,\vec{\mathbf{r}}(M))_q < 0$, we let $r(\cdot)$ and $\hat{r}(\cdot)$ denote the rank functions of M, $M + e$, respectively, and we recall the definition of the matrix entries a_{qS} from (6.20), to justify the following calculation:

$$\begin{aligned}(A\,\vec{\mathbf{r}}(M))_q = \sum_{S\in\mathcal{P}(I)} a_{qS}\, r(S) = \sum_{T\in\mathcal{P}(J)} \alpha_T\, r(q(T)) = \sum_{T\in\mathcal{P}(J)} \alpha_T\, \hat{r}(T)\\ = \vec{\alpha}^{\mathsf{T}}\vec{\mathbf{r}}(M+e) = \Lambda^{\mathsf{T}} B\vec{\mathbf{r}}(M+e) = \Lambda^{\mathsf{T}}\vec{\mathbf{r}}(M) < 0\end{aligned}$$

where the last two steps used Lemma 6.7(iii) and Theorem 6.4 or 6.6 for a field $\mathbb{F}$ of even or odd characteristic, respectively. ■

References

[1] M. Adler, N. J. A. Harvey, K. Jain, R. Kleinberg, and A. R. Lehman, *On the capacity of information networks*, Proc. of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006), pp. 241–250.

[2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, *Network information flow*, IEEE Trans. Inform. Theory **46** (2000), 1204–1216.

[3] N. Alon, A. Hassidim, E. Lubetzky, U. Stav, and A. Weinstein, *Broadcasting with side information*, Proc. of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008), pp. 823–832.

[4] Z. Bar-Yossef, Y. Birk, T. S. Jayram, and T. Kol, *Index coding with side information*, Proc. of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 197–206.

[5] Y. Birk and T. Kol, *Coding-on-demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients*, IEEE Trans. Inform. Theory **52** (2006), 2825–2830. An earlier version appeared in INFOCOM 1998.

[6] R. Dougherty, C. Freiling, and K. Zeger, *Insufficiency of linear coding in network information flow*, IEEE Trans. Inform. Theory **51** (2005), 2745–2759.

[7] R. Dougherty, C. Freiling, and K. Zeger, *Networks, matroids, and non-Shannon information inequalities*, IEEE Trans. Inform. Theory **53** (2007), no. 6, 1949–1969.

[8] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, *XORs in the air: practical wireless network coding*, IEEE/ACM Trans. on Networking **16** (2008), 497–510. An earlier version appeared in SIGCOMM 2006.

[9] E. Lubetzky and U. Stav, *Non-linear index coding outperforming the linear optimum*, IEEE Trans. Inform. Theory **55** (2009), 3544–3551. An earlier version appeared in Proc. of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), pp. 161–167.

[10] S. El Rouayheb, A. Sprintson, and C. Georghiades, *A new construction method for networks from matroids*, IEEE International Symposium on Information Theory (ISIT 2009), pp. 2872–2876.

[11] S. El Rouayheb, A. Sprintson, and C. Georghiades, *On the relation between the Index Coding and the Network Coding problems*, IEEE International Symposium on Information Theory (ISIT 2008), pp. 1823–1827.

[12] R. W. Yeung, S.-Y. R. Li, and N. Cai, *Network coding theory*, Now Publishers Inc, 2006.