Network-Coding Solutions for Minimal Combination Networks and Their Sub-Networks

Han Cai, *Member, IEEE*, Johan Chrisnata, Tuvi Etzion, *Fellow, IEEE*, Moshe Schwartz, *Senior Member, IEEE*, and Antonia Wachter-Zeh, *Member, IEEE*

***Abstract*— Minimal multicast networks are fascinating and efficient combinatorial objects, where the removal of a single link makes it impossible for all receivers to obtain all messages. We study the structure of such networks, and prove some constraints on their possible solutions. We then focus on the combination network, which is one of the simplest and most insightful network in network-coding theory. Of particular interest are minimal combination networks. We study the gap in alphabet size between vector-linear and scalar-linear network-coding solutions for such minimal combination networks and some of their sub-networks. For minimal multicast networks with two source messages we find the maximum possible gap. We define and study sub-networks of the combination network, which we call Kneser networks, and prove that they attain the upper bound on the gap with equality. We also prove that the study of this gap may be limited to the study of sub-networks of minimal combination networks, by using graph homomorphisms connected with the q-analog of Kneser graphs. Additionally, we prove a gap for minimal multicast networks with three or more source messages by studying Kneser networks. Finally, an upper bound on the gap for full minimal combination networks shows nearly no gap, or none in some cases. This is obtained using an MDS-like bound for subspaces over a finite field.**

***Index Terms*— Linear network coding, minimal networks, combination network, graph coloring, q-Kneser graphs.**

I. Introduction

NETWORK coding has been attracting increased attention for almost two decades since the seminal papers [1], [18]. Multicast networks have received most of this attention. A recent survey on the foundations of multicast network coding may be found in [11]. The multicast network-coding problem can be formulated as follows: given an acyclic network with one source which has h messages, for each edge find a function of the messages received at the starting node of the edge, such that each terminal can recover all the messages from its received messages. Such an assignment of a function to each edge is called a *solution* for the network.

Obviously, received messages on an edge can be expressed as functions of the source messages. If these functions are linear, we obtain a *linear solution*. Otherwise, we have a *nonlinear solution*. In linear network coding, each linear function on an edge consists of coding coefficients for each incoming message. If the coding coefficients and the messages are scalars, it is called a *scalar solution*. If the messages are vectors and the coding coefficients are matrices then it is called a *vector solution*. A network which has a solution is called a *solvable* network. It is well known that a multicast network with one source, h messages, and N terminals, is solvable if and only if the min-cut between the source and each terminal is at least h [1], [18].

The minimal *alphabet size*, and in the linear setting, *field size*, of a solution is an important parameter that directly influences the complexity of the calculations at the network nodes. An efficient algorithm to find a field size (not necessarily minimal) that allows a linear solution, and the related linear network code was given in [16]. It is known that any field size $q \geqslant N$ suffices for a linear solution, but it is conjectured that the smallest field size allowing a linear solution is much smaller [10], [11]. This, however, was proved only for two messages [10].

In this work we distinguish between scalar and vector linear solutions. The potential benefits of vector linear solutions have been recognized as early as [21], in which insufficiency of the linear paradigm has been studied in non-multicast networks. We, however, are interested in multicast networks with a linear solution, and take a closer look at the required field size. Given a network $\mathcal{N}$, we define $q_s(\mathcal{N})$ to be the smallest field size q for which $\mathcal{N}$ has a scalar linear solution. Similarly, we define $q_v(\mathcal{N})$ to be the smallest value q^t, q a prime power, such that $\mathcal{N}$ has a vector solution with vectors of length t over a field of size q. By definition,

$$q_s(\mathcal{N}) \geqslant q_v(\mathcal{N}),$$

and we define the *gap* by

$$\mathrm{gap}(\mathcal{N}) \triangleq q_s(\mathcal{N}) - q_v(\mathcal{N}).$$

Manuscript received September 13, 2019; revised February 20, 2020; accepted May 9, 2020. Date of publication May 20, 2020; date of current version October 21, 2020. The work of Tuvi Etzion was supported in part by the Bernard Elkin Chair in Computer Science. The work of Moshe Schwartz and Antonia Wachter-Zeh was supported in part by a German Israeli Project Cooperation (DIP) Grant under Grant PE2398/1-1 and Grant KR3517/9-1. This article was presented at the 2019 International Symposium on Information Theory. *(Corresponding author: Moshe Schwartz.)*

Han Cai and Moshe Schwartz are with the School of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer Sheva 8410501, Israel (e-mail: hancai@aliyun.com; schwartz@ee.bgu.ac.il).

Johan Chrisnata is with the School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 639798, and also with the Department of Computer Science, Technion–Israel Institute of Technology, Haifa 3200003, Israel (e-mail: johan.c@cs.technion.ac.il).

Tuvi Etzion is with the Department of Computer Science, Technion–Israel Institute of Technology, Haifa 3200003, Israel (e-mail: etzion@cs.technion.ac.il).

Antonia Wachter-Zeh is with the Institute for Communications Engineering, Technical University of Munich, 80333 Munich, Germany (e-mail: antonia.wachter-zeh@tum.de).

Communicated by A. Thangaraj, Associate Editor for Coding Techniques.
Digital Object Identifier 10.1109/TIT.2020.2995845

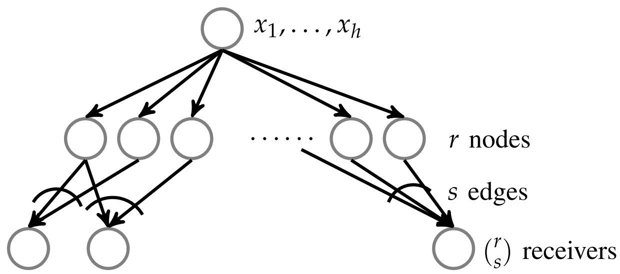

Fig. 1. The $\mathcal{N}_{h,r,s}$ (full) combination network: it has an edge from the source to each of the r nodes in the middle layer. Each of the $\binom{r}{s}$ terminals is connected to a unique set of s middle-layer nodes, and wants to recover all of the h source messages.

While this definition of gap (which was implicitly used in [9]) measures the difference in alphabet size, from an information-theoretic point of view, we might also be interested in the difference in the number of information bits, which we define as

$$\mathrm{gap}_2(\mathcal{N}) \triangleq \log_2 \; q_s(\mathcal{N}) - \log_2 \; q_v(\mathcal{N}).$$

One of the most celebrated families of networks is the family of combination networks [25], which has been used for various topics in network coding, e.g., [13], [14], [20], [23], [28]. The $\mathcal{N}_{h,r,s}$ combination network, where $s \geqslant h$, consists of a single source node with h source messages, r nodes in the middle layer, each s-subset of which is connected to a unique receiver, which is shown in Fig. 1. The network has three layers: the first layer consists of a single source with h messages. The source transmits r messages to the r nodes of the middle layer. Any s nodes in the middle layer are connected to a terminal, and each one of the $\binom{r}{s}$ terminals wants to recover all the h messages. Since we shall also examine sub-networks of the combination network, we will sometimes stress that no part of the network has been removed by saying that the network is a *full* combination network.

It was proved in [29, Chapter 4] that a solution for a combination network exists if and only if a related error-correcting code exists. This network was also generalized to compare scalar and vector network coding [9]. Its sub-networks were used to prove that finding the minimum required field size of a (linear or nonlinear) scalar network code for a certain multicast network is NP-complete [17].

Of particular interest are *minimal multicast networks*. In such a network, the removal of even a single link reduces the cut between the source and at least one of the terminals, which in turn, makes the network unable to transmit all of its messages to all the terminals, i.e., the network becomes unsolvable. In the case of combination networks, only the $\mathcal{N}_{h,r,h}$ networks are minimal. Minimal networks are not only a fascinating extremal combinatorial objects, but also of practical importance since they require the least amount of network resources to enable the multicast operation. Minimal networks have been studied in the past, e.g., [8], [17], but only in the context of scalar network coding.

The goal of this work is to consider two problems which are related to vector coding solutions for minimal combination networks and their sub-networks. Our main contributions are the following: we first study general multicast networks which are minimal. We extend the scalar setting of [8], [17] to the vector setting, and we show minimality entails some structural properties of the graph, as well as some constraints on linear solutions.

We then focus on the gap in general minimal multicast networks. The first to demonstrate a gap exists in some network was [27]. However, the gap there is only 1, and requires at least $h \geqslant 8190$ source messages. This was significantly improved by [9], which showed much larger gaps for as little as $h \geqslant 4$ source messages, in non-minimal networks. In this paper we show that large gaps exist already for $h = 2$ messages, the lowest number of messages for which a gap is possible. Not only that, but the networks we construct are minimal, and we show that they attain the highest possible gap of all minimal multicast networks with two source messages. We further prove that studying the gap in general minimal multicast networks with two source messages is equivalent to studying it only in minimal multicast networks which are sub-networks of combination networks, thus, further motivating the study of combination networks. The main tool we use is a connection with q-Kneser graphs (which are q-analogs of Kneser graphs), as well as a generalization to a new q-Kneser hypergraph.

Finally, we also prove an upper bound on the gap in full minimal combination networks. To the best of our knowledge, this is the first such upper bound on the gap. The bound is obtained using certain longest MDS array codes (e.g., see [4]) or a combinatorial structure which we call an independent configuration. While an upper bound on the length of such codes already exists, we present a different proof approach based on the properties of subspaces in the configuration.

The paper is organized as follows. In Section II we provide the basic notation and definitions used in the paper. Section III studies fundamental properties of general minimal multicast networks. In Section IV we study the gap in networks with two source messages, whereas Section V is devoted to three or more source messages. Section VI focuses on an upper bound on the gap for full minimal combination networks. We conclude in Section VII with a summary of the results and some open problems.

II. Preliminaries

We now provide the basic notation and definitions used throughout the paper. If W is some finite set of n elements, we use $\binom{W}{t}$ to denote the set of all subsets of W of size t. Obviously, using the binomial coefficients,

$$\left|\binom{W}{t}\right| = \binom{n}{t} \triangleq \frac{n!}{t!(n-t)!}.$$

Now, let $\mathbb{F}_q$ denote the finite field of size q, where $q \in \mathbb{P}$ and $\mathbb{P} \subset \mathbb{N}$ denotes the set of prime powers. Taking the q-analog of sets, if V is a vector space over $\mathbb{F}_q$, with $n = \dim V$, we use $\begin{bmatrix} V \\ t \end{bmatrix}$ to denote the set of all vector subspaces of V of dimension t. It is well known that the size of $\begin{bmatrix} V \\ t \end{bmatrix}$ is given by the Gaussian coefficient, namely,

$$\left|\begin{bmatrix} V \\ t \end{bmatrix}\right| = \begin{bmatrix} n \\ t \end{bmatrix}_q \triangleq \frac{\prod_{i=1}^{n}(q^i-1)}{\prod_{i=1}^{t}(q^i-1)\prod_{i=1}^{n-t}(q^i-1)}.$$

We omit the subscript q whenever the field size is understood from the context, and just write $\begin{bmatrix} n \\ t \end{bmatrix}$.

A combinatorial object we shall encounter several times is a spread. If $V = \mathbb{F}_q^n$, a *t-spread*, S, is a collection of subspaces of V of dimension t, i.e., $S \subseteq \begin{bmatrix} V \\ t \end{bmatrix}$, such that the subspaces in S intersect only trivially, and their union is V. Alternatively, every subspace of V of dimension one is contained in exactly one element of S. Thus, the $\begin{bmatrix} n \\ 1 \end{bmatrix}_q$ subspaces of dimension one are equally distributed among the elements of S, and since each t-dimensional subspace contains exactly $\begin{bmatrix} t \\ 1 \end{bmatrix}_q$ subspaces of dimension one, it follows that the number of subspaces in a t-spread is exactly

$$|S| = \frac{\begin{bmatrix} n \\ 1 \end{bmatrix}_q}{\begin{bmatrix} t \\ 1 \end{bmatrix}_q} = \frac{q^n - 1}{q^t - 1}, \tag{1}$$

and it is known that t-spreads exist exactly when $t|n$ (e.g., see [26]).

For any $x \in \mathbb{R}$, $x > 0$ we use $\psi(x)$ to denote the smallest prime power that is greater or equal to x, i.e.,

$$\psi(x) \triangleq \min \{q \in \mathbb{P} \ : \ q \geqslant x\}. \tag{2}$$

By Bertrand's postulate (e.g., see [2]),

$$0 \leqslant \psi(n) - n \leqslant n, \tag{3}$$

for all $n \in \mathbb{N}$. We mention that much stronger results may be obtained at the cost of working only for large enough n. For example, [3] showed that the interval $[x, x + x^{21/40}]$ contains a prime for all large enough x. Thus, for all large enough n,

$$0 \leqslant \psi(n) - n \leqslant n^{21/40}. \tag{4}$$

Our main objects of interest are networks. We shall always assume that our network consists of a finite directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Nodes in the graph will be denoted using lower-case Greek letters. The network contains a single node that is designated as the source, $\sigma \in \mathcal{V}$. It also contains terminal nodes $\mathcal{T} = \{\tau_1, \tau_2, \ldots, \tau_N\} \subseteq \mathcal{V}$. Finally, the source has no incoming edges, but is in possession of h messages, denoted $x_1, \ldots, x_h \in \mathbb{F}_q^t$. We therefore denote this network by $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, h)$.

In the network-coding model, each edge $e \in \mathcal{E}$ carries a value from $\mathbb{F}_q^t$. For a node $\nu \in \mathcal{V}$, let $d_{\text{in}}(\nu)$ denote the in-degree of ν, and let $\text{In}(\nu)$ denote the set of incoming edges into ν, hence $|\text{In}(\nu)| = d_{\text{in}}(\nu)$. For each node ν and each edge e outgoing from ν, the value e carries is a function of the values on edges incoming into ν. In a *linear* setting, this value is a linear combination of the incoming values, where the coefficients are $t \times t$ matrices over $\mathbb{F}_q$.

The goal of the terminals is to gain knowledge of source messages. In the *multicast* setting, every terminal wants to recover all the source messages. When working with vectors from $\mathbb{F}_q^t$, if each of the terminals may recover all the source messages then a *linear solution* to the network is possible, and we say $\mathcal{N}$ has a (q, t)-linear solution. When $t = 1$ we call it a *scalar linear solution*, and in general, a *vector linear solution*.

The source messages naturally form a vector space over $\mathbb{F}_q$ of dimension ht, namely, $\mathbb{F}_q^{ht}$. In the linear case, for each edge $e \in \mathcal{E}$, the values carried by e also form a vector space over $\mathbb{F}_q$, which we denote by $\mathcal{M}(e)$. More precisely, this vector space is given by

$$\mathcal{M}(e) = \left\{ G_e \cdot (x_1 | \ldots | x_h)^T \ : \ x_i \in \mathbb{F}_q^t \right\},$$

where G_e is a $t \times ht$ matrix over $\mathbb{F}_q$. The matrix G_e is called the *global coding matrix* (and when $t = 1$, the *global coding vector*). Obviously,

$$0 \leqslant \dim \mathcal{M}(e) = \text{rank}(G_e) \leqslant t.$$

We also say each vertex $\nu \in \mathcal{V}$ has access to

$$\mathcal{M}(\nu) \triangleq \left\{ G_\nu \cdot (x_1 | \ldots | x_h)^T \ : \ x_i \in \mathbb{F}_q^t \right\}, \tag{5}$$

where G_ν is the $(d_{\text{in}}(\nu) \cdot t) \times ht$ matrix over $\mathbb{F}_q$ defined by

$$G_\nu \triangleq \begin{pmatrix} G_{e_1} \\ G_{e_2} \\ \vdots \\ G_{e_{d_{\text{in}}(\nu)}} \end{pmatrix},$$

with $\text{In}(\nu) = \left\{ e_1, \ldots, e_{d_{\text{in}}(\nu)} \right\}$. Thus,

$$\dim \mathcal{M}(\nu) = \text{rank}(G_\nu) \leqslant \min(d_{\text{in}}(\nu) \cdot t, ht). \tag{6}$$

Combining these facts together, a terminal $\tau \in \mathcal{T}$ is successful in the linear multicast setting if and only if

$$\dim \mathcal{M}(\tau) = \text{rank}(G_\tau) = ht.$$

A node in the network is said to be *essential* if it is on a path from the source to some terminal. We assume throughout the paper that all nodes are essential, since non-essential nodes may be removed without affecting the solution.

A *cut* in the network is defined by a partition $\mathcal{V} = S \cup T$, with the source $\sigma \in S$, and some terminal $\tau_i \in T$. We say the size of the cut is m if there are exactly m edges crossing it from S to T. It is well known (see [11]) that in a multicast network there exists a (q, t)-linear solution if and only if every cut has size at least h.

Given a network $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, h)$, we define $q_s(\mathcal{N})$ to be the smallest field size q for which $\mathcal{N}$ has a $(q, 1)$-scalar linear solution. Such a solution is said to be *scalar-optimal*. We also define $q_v(\mathcal{N})$ to be the smallest value q^t such that $\mathcal{N}$ has a (q, t)-linear solution for some t, and such a solution is said to be *vector-optimal*. By definition, $q_s(\mathcal{N}) \geqslant q_v(\mathcal{N})$, and so we define the *(vector) gap* by

$$\text{gap}(\mathcal{N}) \triangleq q_s(\mathcal{N}) - q_v(\mathcal{N}),$$

and the *information gap* by

$$\text{gap}_2(\mathcal{N}) \triangleq \log_2 \ q_s(\mathcal{N}) - \log_2 \ q_v(\mathcal{N}).$$

Finally, as mentioned in the previous section, a celebrated family of networks that has been studied extensively is the family of combination networks. We briefly recall its definition: for $h, r, s \in \mathbb{N}$, $r \geqslant s$, the $\mathcal{N}_{h,r,s}$ combination network (depicted in Fig. 1) is a multicast network with a single source σ, connected to r nodes in the middle layer, which we denote $\mathcal{L} = \{\lambda_1, \ldots, \lambda_r\}$. We then have $\binom{r}{s}$ terminal nodes, which we may think of as indexed by $\binom{\mathcal{L}}{s}$. Each of the terminals is connected to a unique subset of s nodes from the middle layer. These networks and their sub-networks will be the focus of this work.

III. Minimal Multicast Networks

A *minimal* multicast network can deliver h messages from the source to each of the terminals while each of its proper sub-networks (containing all the original terminals) can deliver at most $h-1$ messages to at least one of the terminals. From a practical point of view, considering such minimal networks is interesting as it minimizes the used network resources. From a theoretical perspective, minimal networks are a fascinating extremal combinatorial object. Minimal networks have been considered in the past [8], [17], however only in the scalar case. The results of this section may be considered as a generalization of these works.

Definition 1: A multicast network $\mathcal{N}=(\mathcal{G},\sigma,\mathcal{T},h)$ is said to be *minimal* if every edge crosses a cut of size h.

Thus, in a minimal network, the removal of any edge from $\mathcal{E}$ makes at least one cut have size strictly less than h, and therefore unsolvable.

Given a directed graph $\mathcal{G}=(\mathcal{V},\mathcal{E})$, and $m\in\mathbb{N}$, we define the m-parallelized version of it as $m\mathcal{G}=(\mathcal{V},m\mathcal{E})$, where $m\mathcal{E}$ denotes the multiset obtained from $\mathcal{E}$ by having each element appear m times. In essence, in the m-parallelized graph we keep the same nodes, but duplicate each edge m times.

Lemma 2: If $\mathcal{N}=(\mathcal{G},\sigma,\mathcal{T},h)$ is a minimal multicast network with a (q,t)-linear solution, then $\mathcal{N}'=(t\mathcal{G},\sigma,\mathcal{T},ht)$ is a minimal multicast network with a $(q,1)$-scalar linear solution.

Proof: The claim is entirely trivial: each transmitted vector on an edge $\nu_1\xrightarrow{e}\nu_2$ in $\mathcal{N}$ is broken up into its t components which are then transmitted separately over the corresponding t parallel edges in $\mathcal{N}'$. These are clearly linear combinations (with scalar coefficients from $\mathbb{F}_q$) of the values entering the node ν_1 in $\mathcal{N}'$. Finally, the size of each cut in $\mathcal{N}'$ is obviously t times the size of the same cut in $\mathcal{N}$, hence the minimality of $\mathcal{N}'$. ■

Lemma 3: Consider a minimal multicast network $\mathcal{N}=(\mathcal{G},\sigma,\mathcal{T},h)$ with a (q,t)-linear solution. Then

$$\dim\mathcal{M}(\nu)=d_{\mathrm{in}}(\nu)\cdot t$$

for each $\nu\in\mathcal{V}\setminus\{\sigma\}$, where $\mathcal{M}(\nu)$ was defined in (5).

Proof: Define $\mathcal{N}'=(t\mathcal{G},\sigma,\mathcal{T},ht)$, and construct a $(q,1)$-scalar linear solution to $\mathcal{N}'$ from the solution to $\mathcal{N}$, as described in the proof of Lemma 2. It is obvious that $\mathcal{M}(\nu)$ remains unchanged for all $\nu\in\mathcal{V}$. To avoid confusion, we let $d_{\mathrm{in}}^{\mathcal{G}}$ denote the in-degree in the graph $\mathcal{G}$, and $d_{\mathrm{in}}^{t\mathcal{G}}$ the in-degree in the t-parallelized graph $t\mathcal{G}$. Assume to the contrary that $\dim\mathcal{M}(\nu)<t\cdot d_{\mathrm{in}}^{\mathcal{G}}(\nu)=d_{\mathrm{in}}^{t\mathcal{G}}(\nu)$. Then, in the network $\mathcal{N}'$ there exists an edge $e\in\mathrm{In}(\nu)$ that always carries some fixed linear combination of other edges in $\mathrm{In}(\nu)$. Removing e still allows a solution to $\mathcal{N}'$ since the original $\mathcal{M}(\nu)$ may still be computed from the surviving edges, and therefore the original values on edges leaving ν may be computed as a linear combination as well. However, this contradicts the minimality of $\mathcal{N}'$ obtained by Lemma 2. ■

The following corollary is trivial in the scalar regime, appearing as a side note in [8]. Using the previous lemmas it may also be proved for the general case.

Corollary 4: In a minimal multicast network $\mathcal{N}=(\mathcal{G},\sigma,\mathcal{T},h)$ with a (q,t)-linear solution we have $d_{\mathrm{in}}(\nu)\leqslant h$ for all $\nu\in\mathcal{V}\setminus\{\sigma\}$.

Proof: Combining Lemma 3 with (6) we get

$$d_{\mathrm{in}}(\nu)\cdot t=\dim\mathcal{M}(\nu)\leqslant ht.$$

The claim follows immediately. ■

IV. Minimal Networks With Two Source Messages

In this section we focus on the case of two source messages, i.e., $h=2$. Our goal is to show that there exist networks with two source messages with a positive gap. Such networks have not been demonstrated in the past, and a gap was shown to exist only in networks with at least three messages [9]. A part of the method we describe here bears some resemblance to the ones described in [8], [10], [17]. The resulting networks are in fact minimal, as well as sub-networks of the (minimal) combination network $\mathcal{N}_{h,r,h}$. Additionally, we show that the networks we construct attain the largest possible gap of any minimal multicast network with two source messages.

Of particular interest in what follows, will be the q-analog of the Kneser graph, denoted $qK_{n:m}$, whose set of vertices is $\begin{bmatrix}V\\m\end{bmatrix}$, where $V=\mathbb{F}_q^n$, and an undirected edge connects two vertices iff the corresponding m-dimensional subspaces have a trivial intersection (e.g., see [6], [7], [15] and references therein).

We also recall graph homomorphisms, which act as a generalization of graph coloring. Given two undirected graphs, $\mathcal{G}_1=(\mathcal{V}_1,\mathcal{E}_1)$ and $\mathcal{G}_2=(\mathcal{V}_2,\mathcal{E}_2)$, we say $\phi:\mathcal{V}_1\to\mathcal{V}_2$ is a *homomorphism* if $\{\nu,\nu'\}\in\mathcal{E}_1$ implies $\{\phi(\nu),\phi(\nu')\}\in\mathcal{E}_2$ for all $\nu,\nu'\in\mathcal{E}_1$. As is customary, we shall denote this homomorphism simply by writing $\mathcal{G}_1\to\mathcal{G}_2$. This is a generalization of coloring since a homomorphism $\mathcal{G}_2\to K_n$ (where K_n denotes the complete graph on n vertices) is equivalent to a coloring of $\mathcal{G}_2$ with n colors (think of each of the vertices of K_n representing a distinct color, which is assigned to the vertices in the reverse image of the homomorphism). Additionally, it is easy to see that $\chi(\mathcal{G}_2)$ (the chromatic number of $\mathcal{G}_2$) is the minimum n such that there exists a homomorphism $\mathcal{G}_2\to K_n$. Since homomorphisms are easily seen to be transitive, i.e., $\mathcal{G}_1\to\mathcal{G}_2\to\mathcal{G}_3$ implies $\mathcal{G}_1\to\mathcal{G}_3$, we therefore have that $\mathcal{G}_1\to\mathcal{G}_2$ implies $\chi(\mathcal{G}_1)\leqslant\chi(\mathcal{G}_2)$.

We now continue with the case of two source messages. Let $\mathcal{G}=(\mathcal{V},\mathcal{E})$ be a finite directed acyclic graph. We describe the following construction of an undirected graph, $\mathrm{skel}(\mathcal{G})$, which we call the *skeleton of* $\mathcal{G}$.

Construction A: Let $\mathcal{G}=(\mathcal{V},\mathcal{E})$ be a finite directed acyclic graph. Define

$$\mathcal{V}_{\neq1}\triangleq\{\nu\in\mathcal{V}\ :\ d_{\mathrm{in}}(\nu)\neq1\},$$
$$\mathcal{E}_{\neq1}\triangleq\{(\nu,\nu')\in\mathcal{E}\ :\ \nu\in\mathcal{V}_{\neq1},\nu'\in\mathcal{V}\}.$$

Additionally, for each $e\in\mathcal{E}_{\neq1}$ we define $T(e)$ to be the set of all edges $e'\in\mathcal{E}$ that may be reached from e by directed paths in $\mathcal{G}$ that pass only through vertices in $\mathcal{V}\setminus\mathcal{V}_{\neq1}$. Namely, for an edge $e=(v,u)$ with $d_{\mathrm{in}}(v)>1$, $T(e)$ is the set of edges reachable from e along paths that pass only through vertices with in-degree one. The set $T(e)$ describes a tree with root v.

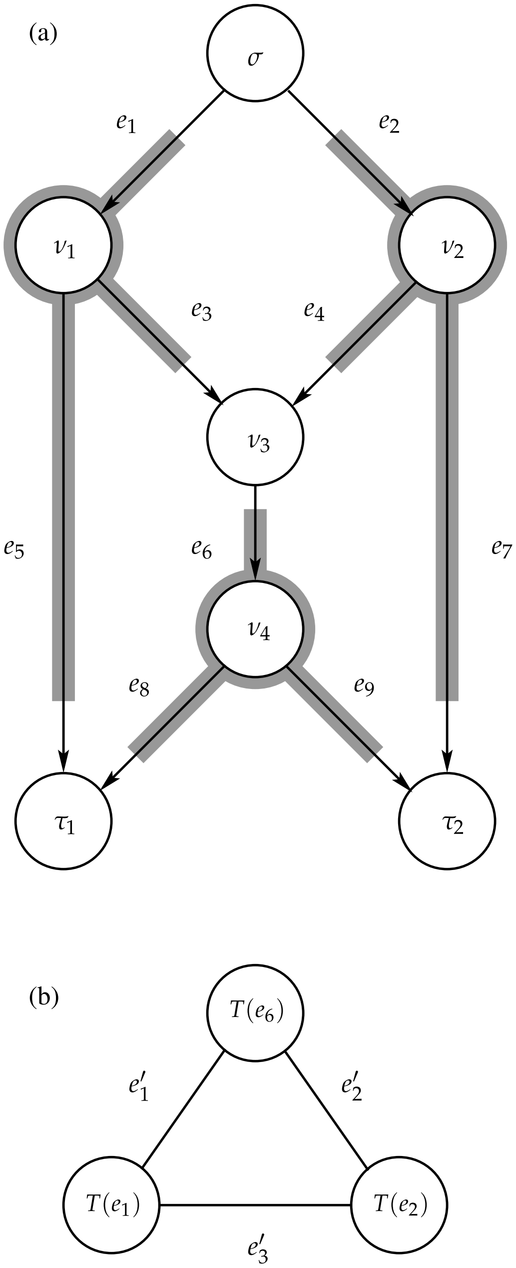

Fig. 2. (a) The famous butterfly network $\mathcal{N} = (\mathcal{G}, \sigma, \{\tau_1, \tau_2\}, 2)$, and (b) its skeleton $\text{skel}(\mathcal{G})$. In the upper figure, the shaded parts show the trees $T(e_1)$, $T(e_2)$, and $T(e_6)$.

We now define the skeleton graph $\text{skel}(\mathcal{G}) = (\mathcal{V}', \mathcal{E}')$. The vertex set is defined as

$$\mathcal{V}' = \{T(e) \ : \ e \in \mathcal{E}_{\neq 1}\},$$

namely, a node for each tree $T(e)$. The edge set is defined by

$$\mathcal{E}' = \{\{T(e_1), T(e_2)\} \ : \ T(e_1), T(e_2) \in \mathcal{V}', T(e_1) \neq T(e_2), \exists (\nu_1, \nu), (\nu_2, \nu) \in \mathcal{E} \text{ s.t. } (\nu_1, \nu) \in T(e_1), (\nu_2, \nu) \in T(e_2)\},$$

namely, two trees, $T(e_1)$ and $T(e_2)$, induce an edge between their respective nodes in $\text{skel}(\mathcal{G})$, if an edge in $T(e_1)$ and an edge in $T(e_2)$ enter the same node in $\mathcal{G}$. □

Example 5: Figure 2(a) shows the famous butterfly network, $\mathcal{N} = (\mathcal{G}, \sigma, \{\tau_1, \tau_2\}, 2)$. For the graph $\mathcal{G}$ we have (in terms of Construction A) that

$$\mathcal{V}_{\neq 1} = \{\sigma, \nu_3, \tau_1, \tau_2\}.$$

Note that we must always have $\sigma \in \mathcal{V}_{\neq 1}$ since $d_{\text{in}}(\sigma) = 0$. Additionally,

$$\mathcal{E}_{\neq 1} = \{e_1, e_2, e_6\},$$

i.e., the set of edges whose source has in-degree that is not 1. We also have

$$T(e_1) = \{e_1, e_3, e_5\},$$
$$T(e_2) = \{e_2, e_4, e_7\},$$
$$T(e_6) = \{e_6, e_8, e_9\}.$$

Each of these sets becomes a vertex in $\text{skel}(\mathcal{G})$, shown in Figure 2(b).

In [17] only the reverse process, i.e., mapping what we call $\text{skel}(\mathcal{G})$ back into a network, is used. This is later also cited and used in [8]. The forward process of mapping a network to its skeleton bears some resemblance to the procedure described in [10]. There, however, a variation on the line graph is used and not the graph. Additionally, the tree decomposition described in [10] creates a directed graph unlike the undirected skeleton we have here. The trees in the decomposition of [10] form a subset (sometimes a proper subset) of $\{T(e) \ : \ e \in \mathcal{E}_{\neq 1}\}$ which was defined here.

A simple observation is proved in the next lemma.

Lemma 6: In Construction A, the set $\mathcal{P} \triangleq \{T(e) \ : \ e \in \mathcal{E}_{\neq 1}\}$ is a partition of $\mathcal{E}$.

Proof: Since $\mathcal{G}$ is directed and acyclic, find a topological ordering of its vertices. Assume to the contrary that $\mathcal{P}$ is not a partition of $\mathcal{E}$, which implies an edge $e_1 = (\nu_1, \nu_2) \in \mathcal{E}$ is not a member of any of the sets in $\mathcal{P}$ or a member of two sets in $\mathcal{P}$. Assume further, without loss of generality, that among all such edges, e_1 denotes an edge for which ν_1 is minimal under the topological ordering.

If $\nu_1 \in \mathcal{V}_{\neq 1}$, then by construction, $e_1 \in T(e_1) \in \mathcal{P}$ uniquely. Otherwise, there exists exactly a single edge $e_0 = (\nu_0, \nu_1) \in \mathcal{E}$. By the minimality of our choice of e_1, we have $e_0 \in T(e) \in \mathcal{P}$ uniquely, for some $e \in \mathcal{E}_{\neq 1}$. By construction, $e_1 \in T(e) \in \mathcal{P}$ uniquely as well. ■

The following lemma connects the existence of a (q, t)-linear solution to the existence of a certain graph homomorphism when we have $h = 2$ source messages.

Lemma 7: Let $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$ be a minimal multicast network. Then $\mathcal{N}$ has a (q, t)-linear solution if and only if $\text{skel}(\mathcal{G}) \to qK_{2t:t}$, i.e., there is a homomorphism from $\text{skel}(\mathcal{G})$ to $qK_{2t:t}$.

Proof: For the only if part, assume $\mathcal{N}$ has a (q, t)-linear solution. We assume without loss of generality that nodes with in-degree 1 simply forward their incoming message. Thus, given any $e \in \mathcal{E}_{\neq 1}$, all the edges $e' \in T(e)$ carry the same message, i.e., $\mathcal{M}(e') = \mathcal{M}(e)$. We construct the homomorphism $\phi : \text{skel}(\mathcal{G}) \to qK_{2t:t}$ by setting $\phi(T(e)) = \mathcal{M}(e)$ for each $e \in \mathcal{E}_{\neq 1}$.

We verify this is indeed a homomorphism, since an edge $\{T(e_1), T(e_2)\}$ in $\text{skel}(\mathcal{G})$ corresponds to edges $(\nu_1, \nu) \in T(e_1)$ and $(\nu_2, \nu) \in T(e_2)$ in $\mathcal{G}$. By Corollary 4, $d_{\text{in}}(\nu) = 2$. By Lemma 3, $\dim \mathcal{M}(\nu) = 2t$, hence $\dim \mathcal{M}(e_1) =$

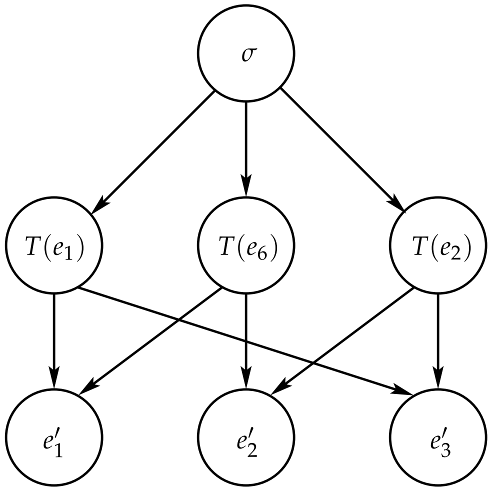

Fig. 3. The sub-network of a combination network that is constructed in Lemma 8, which has the same skeleton as the butterfly network of Figure 2.

$\dim \mathcal{M}(e_2) = t$ and $\mathcal{M}(e_1), \mathcal{M}(e_2)$ must intersect trivially. Thus, $\mathcal{M}(e_1)$ and $\mathcal{M}(e_2)$ are joined by an edge in $qK_{2t:t}$.

For the if part, assume we have a homomorphism $\phi : \mathrm{skel}(\mathcal{G}) \to qK_{2t:t}$. It is easy to verify that setting $\mathcal{M}(e') = \phi(T(e))$, for all $e' \in T(e)$, $e \in \mathcal{E}_{\neq 1}$, is indeed a (q,t)-linear solution, since nodes with in-degree 1 simply forward their incoming message, whereas nodes with in-degree two receive two trivially intersecting subspaces, and may therefore reconstruct the entire space $\mathbb{F}_q^{2t}$. Such nodes can linearly create any subspace as an outgoing message. ∎

We now prove that the mapping $\mathrm{skel}(\cdot)$ of minimal multicast networks with two messages to undirected graphs, is surjective. This already appears in part in [17].

Lemma 8: Let $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ be a finite undirected graph with no isolated vertices. Then there exists a minimal multicast network $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$ such that $\mathrm{skel}(\mathcal{G}) = \mathcal{G}'$.

Proof: Construct the desired network $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$ as follows. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and fix $\mathcal{V} = \mathcal{V}' \cup \mathcal{E}' \cup \{\sigma\}$. Additionally, the directed edges are defined as

$$\mathcal{E} = \{(\sigma, \nu) \ : \ \nu \in \mathcal{V}'\} \cup \bigcup_{e=\{\nu,\nu'\}\in\mathcal{E}'} \{(\nu, e), (\nu', e)\}.$$

The source is σ, and the terminals are $\mathcal{T} = \mathcal{E}'$.

Intuitively, we construct the network by adding a new source node σ, which is connected to each of the original nodes of $\mathcal{G}'$. Every edge $e = \{\nu, \nu'\}$ in the original graph $\mathcal{G}'$ defines a new terminal node into which an edge from ν and an edge from ν' are connected. It is easy to verify that $\mathrm{skel}(\mathcal{G}) = \mathcal{G}'$, and that $\mathcal{N}$ is minimal. The requirement that there exist no isolated vertices ensures all nodes in the middle layer are essential. ∎

Example 9: Continuing Example 5, in this example we use Lemma 8 to construct a sub-network of a combination network whose skeleton is also $\mathrm{skel}(\mathcal{G})$, where $\mathcal{G}$ is the butterfly network from Example 5. Recall that $\mathrm{skel}(\mathcal{G})$ is depicted in Figure 2(b). The construction is done by adding a source node σ (with h source messages and no incoming edges), each vertex in $\mathrm{skel}(\mathcal{G})$ becomes a node in the middle layer of the constructed graph, and each edge in $\mathrm{skel}(\mathcal{G})$ becomes a terminal node. The result is depicted in Figure 3.

We note that the network constructed in Lemma 8 is in fact a sub-network of the combination network $\mathcal{N}_{2,r,2}$. Thus, as a corollary we obtain that the study of minimal multicast networks with two source messages may be restricted to sub-networks of combination networks of the form $\mathcal{N}_{2,r,2}$.

Corollary 10: For any minimal multicast network $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$ there exists a minimal multicast network $\mathcal{N}' = (\mathcal{G}', \sigma', \mathcal{T}', 2)$ such that $\mathcal{G}'$ is a subgraph of a combination network, as well as $\mathrm{gap}(\mathcal{N}) = \mathrm{gap}(\mathcal{N}')$ and $\mathrm{gap}_2(\mathcal{N}) = \mathrm{gap}_2(\mathcal{N}')$.

Proof: We first note that, by construction, $\mathrm{skel}(\mathcal{G})$ does not have isolated nodes as these would imply non-essential nodes in $\mathcal{G}$. We then use Lemma 8 to construct $\mathcal{N}'$ such that $\mathrm{skel}(\mathcal{G}') = \mathrm{skel}(\mathcal{G})$. By Lemma 7, the networks $\mathcal{N}$ and $\mathcal{N}'$ have (q,t)-linear solutions for exactly the same pairs $(q,t) \in \mathbb{P} \times \mathbb{N}$. Thus, the gaps are the same. Finally, as observed, $\mathcal{N}'$ is a sub-network of a minimal combination network $\mathcal{N}_{2,r,2}$. ∎

We can now restate the definition of the gap for minimal multicast networks $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$. First, we define $\mathcal{H} = \mathrm{skel}(\mathcal{G})$. Then, by Lemma 7,

$$q_v(\mathcal{N}) = \min \left\{ q^t \ : \ q \in \mathbb{P}, t \in \mathbb{N}, \exists \mathcal{H} \to qK_{2t:t} \right\},$$
$$q_s(\mathcal{N}) = \min \left\{ q \ : \ q \in \mathbb{P}, \exists \mathcal{H} \to qK_{2:1} \right\}.$$

We now crucially observe that $qK_{2:1} \cong K_{q+1}$. To see this we first note that the number of vertices in $qK_{2:1}$ is

$$\begin{bmatrix} 2 \\ 1 \end{bmatrix}_q = \frac{q^2 - 1}{q - 1} = q + 1.$$

We then note that any two distinct 1-dimensional subspaces of $\mathbb{F}_q^2$ have trivial intersection, and hence, any two distinct vertices in $qK_{2:1}$ are connected by an edge. Thus, the graph $qK_{2:1}$ is a complete graph on $q+1$ vertices, which is denoted as K_{q+1}. We therefore have

$$q_s(\mathcal{N}) = \psi(\chi(\mathcal{H}) - 1),$$

where we use ψ from (2), since $q_s(\mathcal{N})$ is required to be in $\mathbb{P}$.

We recall some known results on the chromatic number of certain q-Kneser graphs.

Theorem 11 [6], [7], [15]: Let $q \in \mathbb{P}$ and $t \in \mathbb{N}$. Then,

$$\chi(qK_{2t:t}) \leqslant q^t + q^{t-1}.$$

If $q \geqslant 5$ or $t \leqslant 3$, then

$$\chi(qK_{2t:t}) = q^t + q^{t-1}.$$

We also prove the following lemma which shall become useful soon.

Lemma 12: Let $q, \overline{q} \in \mathbb{P}$ and $t, \overline{t} \in \mathbb{N}$. If $q^t > \overline{q}^{\overline{t}}$ then

$$qK_{2t:t} \not\to \overline{q}K_{2\overline{t}:\overline{t}}.$$

Proof: An upper bound on the size of any clique in $qK_{2t:t}$ is given by $(q^{2t} - 1)/(q^t - 1) = q^t + 1$, describing a scenario where (except for the zero vector) the subspaces corresponding to the nodes in the clique form a partition of $\mathbb{F}_q^{2t}$ (except for the zero vector). This scenario is indeed always possible, and is attained by a t-spread of $\mathbb{F}_q^{2t}$, and the upper bound follows from (1).

Assume to the contrary that there exists a homomorphism $qK_{2t:t} \to \overline{q}K_{2\overline{t}:\overline{t}}$. Since q-Kneser graphs never contain self loops, the nodes in a clique in $qK_{2t:t}$ are mapped to distinct nodes of a clique in $\overline{q}K_{2\overline{t}:\overline{t}}$. But by assumption, $q^t+1 > \overline{q}^{\overline{t}}+1$, namely, the size of largest clique in $qK_{2t:t}$ is strictly larger than the size of the largest clique in $\overline{q}K_{2\overline{t}:\overline{t}}$, a contradiction. ∎

We are now in a position to provide minimal multicast networks with two messages with a positive gap. The key to proving this claim is to build networks whose skeleton is a q-Kneser graph. We define these networks in more generality than required here, since they will be used in the following section as well.

Definition 13: Let $q \in \mathbb{P}$, and $h, t \in \mathbb{N}$, $h \geqslant 2$. We define the *Kneser network*, $\mathcal{K}_{q,t;h}$ in the following way: the network has a single source node σ. Denote $V \triangleq \mathbb{F}_q^{ht}$. The source node is connected to $\begin{bmatrix} ht \\ t \end{bmatrix}$ nodes in the middle layer, which are indexed by $\begin{bmatrix} V \\ t \end{bmatrix} = \{V_1, V_2, \dots\}$. Finally, if we have indices $1 \leqslant i_1 < i_2 < \cdots < i_h \leqslant \begin{bmatrix} ht \\ t \end{bmatrix}$ such that

$$V_{i_1} + \cdots + V_{i_h} = V = \mathbb{F}_q^{ht}, \tag{7}$$

then their corresponding nodes in the middle layer are connected to a unique terminal node.

We observe that for $h = 2$,

$$\mathrm{skel}(\mathcal{K}_{q,t;2}) = qK_{2t:t}. \tag{8}$$

We can therefore determine its q_v exactly, which we do in the following lemma.

Lemma 14: For all $q \in \mathbb{P}$, $t \in \mathbb{N}$,

$$q_v(\mathcal{K}_{q,t;2}) = q^t.$$

Proof: By (8), and since the identity homomorphism $qK_{2t:t} \to qK_{2t:t}$ always exists, we have

$$q_v(\mathcal{K}_{q,t;2}) \leqslant q^t.$$

By Lemma 12, this becomes an equality, as claimed. ∎

We now show a gap exists in many Kneser networks with two source messages.

Theorem 15: For all $q \in \mathbb{P}$ and $t \in \mathbb{N}$, with $q \geqslant 5$ or $t \leqslant 3$,

$$\begin{aligned} \mathrm{gap}(\mathcal{K}_{q,t;2}) &= \psi(q^t + q^{t-1} - 1) - q^t \geqslant q^{t-1} - 1, \\ \mathrm{gap}_2(\mathcal{K}_{q,t;2}) &= \log_2 \psi(q^t + q^{t-1} - 1) - \log_2\ q^t \\ &\geqslant \log_2 \left(1 + \frac{1}{q} - \frac{1}{q^t}\right). \end{aligned}$$

Proof: By Lemma 14,

$$q_v(\mathcal{K}_{q,t;2}) = q^t.$$

However, by Theorem 11, the chromatic number of $qK_{2t:t}$ is

$$\chi(qK_{2t:t}) = q^t + q^{t-1},$$

for $q \geqslant 5$ or $t \leqslant 3$. Thus, for these cases

$$q_s(\mathcal{K}_{q,t;2}) = \psi(q^t + q^{t-1} - 1),$$

and the equality part of the claim is proved. The claimed inequality follows from $\psi(n) \geqslant n$ for all $n \in \mathbb{N}$. ∎

This result is matched by the following upper bound.

Theorem 16: If $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$ is a minimal multicast network with a (q,t)-vector-optimal linear solution, then

$$\begin{aligned} \mathrm{gap}(\mathcal{N}) &\leqslant \psi(q^t + q^{t-1} - 1) - q^t \leqslant q^t + 2q^{t-1} - 2, \\ \mathrm{gap}_2(\mathcal{N}) &\leqslant \log_2 \psi(q^t + q^{t-1} - 1) - \log_2\ q^t \\ &\leqslant 1 + \log_2 \left(1 + \frac{1}{q} - \frac{1}{q^t}\right). \end{aligned}$$

Proof: By definition, a (q,t)-vector-optimal linear solution implies $q_v(\mathcal{N}) = q^t$, and we have

$$\mathrm{skel}(\mathcal{G}) \to qK_{2t:t}.$$

By Theorem 11 we know that $\chi(qK_{2t:t}) \leqslant q^t + q^{t-1}$. We denote $q' = \psi(q^t + q^{t-1} - 1)$, and therefore

$$qK_{2t:t} \to K_{\chi(qK_{2t:t})} \to K_{q^t+q^{t-1}} \to K_{q'+1} \cong q'K_{2:1},$$

where the third homomorphism follows trivially from the fact that $q'+1 \geqslant q^t + q^{t-1}$. By the transitivity of homomorphisms we have

$$\mathrm{skel}(\mathcal{G}) \to q'K_{2:1},$$

and then by Lemma 7, the network $\mathcal{N}$ has a $(q',1)$-linear solution, namely,

$$q_s(\mathcal{N}) \leqslant q' = \psi(q^t + q^{t-1} - 1).$$

Finally, the last inequality follows from (3). ∎

Corollary 17: Kneser networks with two source messages, $\mathcal{K}_{q,t;2}$, $q \in \mathbb{P}$, $t \in \mathbb{N}$, and $q \geqslant 5$ or $t \leqslant 3$, attain the maximum possible gap among all minimal networks with two source messages and (q,t)-vector-optimal linear solutions.

A few cases remain uncovered by our previous treatment, namely, $q \in \{2,3,4\}$ and $t \geqslant 4$. We can show a gap for these cases as well, albeit a much smaller guaranteed gap of merely 1. To that end we recall some known facts about the q-analog of the renowned Erdős-Ko-Rado Theorem. If $V \triangleq \mathbb{F}_q^n$, then $\mathcal{F} \subseteq \begin{bmatrix} V \\ \ell \end{bmatrix}$ is an *m-intersecting family* if for all $W, W' \in \mathcal{F}$, $\dim(W \cap W') \geqslant m$. The maximal size of an m-intersecting family was found in [12]:

Theorem 18 [12]: Let $n \geqslant 2\ell - m$, $V \triangleq \mathbb{F}_q^n$, and let $\mathcal{F} \subseteq \begin{bmatrix} V \\ \ell \end{bmatrix}$ be an m-intersecting family. Then

$$|\mathcal{F}| \leqslant \max\left\{ \begin{bmatrix} n-m \\ \ell-m \end{bmatrix}, \begin{bmatrix} 2\ell-m \\ \ell \end{bmatrix} \right\}.$$

If $2\ell - m < n < 2\ell$, this is attained uniquely by,

$$\mathcal{F} = \left\{ U \in \begin{bmatrix} V \\ \ell \end{bmatrix} \ :\ U \subseteq V_{2\ell-m} \right\}, \tag{9}$$

where $V_{2\ell-m} \in \begin{bmatrix} V \\ 2\ell-m \end{bmatrix}$ is arbitrary. If $n > 2\ell$, this is attained uniquely by,

$$\mathcal{F} = \left\{ U \in \begin{bmatrix} V \\ \ell \end{bmatrix} \ :\ V_m \subseteq U \right\}, \tag{10}$$

where $V_m \in \begin{bmatrix} V \\ m \end{bmatrix}$ is arbitrary. If $n = 2\ell$ then both (9) and (10) attain the maximum, and if additionally $m = 1$, it is uniquely so.

In the context of coding over subspaces, m-intersecting families have been studied as anticodes. If $V \triangleq \mathbb{F}_q^n$, the distance between $W, W' \in \begin{bmatrix} V \\ \ell \end{bmatrix}$ is defined as $d(W, W') \triangleq$

$\ell - \dim(W \cap W')$. An anticode of diameter D is a set $\mathcal{F} \subseteq \begin{bmatrix} V \\ \ell \end{bmatrix}$ such that $W, W' \in \mathcal{F}$ implies $d(W, W') \leqslant D$. Thus, an m-intersecting family is equivalent to an anticode of diameter $\ell - m$. Of particular interest to us is the following theorem from [26] (stated originally in the anticode terminology):

Theorem 19 [26]: Let $n = 2\ell$, $1 \leqslant m \leqslant \ell$, and $V \triangleq \mathbb{F}_q^n$. If $\mathcal{F}_1, \dots, \mathcal{F}_c \subseteq \begin{bmatrix} V \\ \ell \end{bmatrix}$, and each $\mathcal{F}_i$ is an m-intersecting family of type (9) or (10), then they cannot form a partition (tiling) of $\begin{bmatrix} V \\ \ell \end{bmatrix}$.

We are now in a position to show the gap.

Theorem 20: For all $q \in \mathbb{P}$ and all $t \in \mathbb{N}$, $t \geqslant 2$,

$$\operatorname{gap}(\mathcal{K}_{q,t;2}) \geqslant \psi(q^t+1) - q^t \geqslant 1,$$
$$\operatorname{gap}_2(\mathcal{K}_{q,t;2}) \geqslant \log_2 \psi(q^t+1) - \log_2\ q^t \geqslant \log_2\left(1 + \frac{1}{q^t}\right).$$

Proof: Again, by Lemma 14, $q_v(\mathcal{K}_{q,t;2}) = q^t$. We now contend that $\chi(qK_{2t:t}) > q^t + 1$. Recall that the vertex set of $qK_{2t:t}$ is $\begin{bmatrix} V \\ t \end{bmatrix}$, where $V = \mathbb{F}_q^{2t}$. Assume a coloring of $qK_{2t:t}$ with c colors. Let $U_i \subset \begin{bmatrix} V \\ t \end{bmatrix}$, $1 \leqslant i \leqslant c$, be the set of vertices colored with color i. Then each U_i is a 1-intersecting family. Also, the set $\{U_i\}_{1 \leqslant i \leqslant c}$ forms a partition of $\begin{bmatrix} V \\ t \end{bmatrix}$.

By Theorem 18, for all $1 \leqslant i \leqslant c$,

$$|U_i| \leqslant \begin{bmatrix} 2t-1 \\ t-1 \end{bmatrix},$$

and U_i is either of type (9) or (10) if equality holds. However, by Theorem 19, there is no tiling of $\begin{bmatrix} V \\ t \end{bmatrix}$ by U_i if they are of maximum size. Thus,

$$\chi(qK_{2t:t}) > \frac{\begin{bmatrix} 2t \\ t \end{bmatrix}}{\begin{bmatrix} 2t-1 \\ t-1 \end{bmatrix}} = q^t + 1.$$

It follows that

$$\operatorname{gap}(\mathcal{K}_{q,t;2}) = q_s(\mathcal{K}_{q,t;2}) - q_v(\mathcal{K}_{q,t;2}) \geqslant \psi(q^t+1) - q^t \geqslant 1,$$

and similarly for $\operatorname{gap}_2(\mathcal{K}_{q,t;2})$, as claimed. ■

V. Minimal Networks With More Than Two Source Messages

Prior to this work, a gap was shown to exist only in networks with at least four source messages [9], and the networks there are not minimal. An additional ad-hoc example with three source messages was also given in [9] but not in the form of a gap. Instead, it was shown that for the same field size, a different number of nodes in the middle layer of the network is required.

In the previous section we showed a gap exists for $h = 2$ source messages. We first contend this immediately shows a gap exists in networks with any $h \geqslant 3$ source messages. We show this by giving a general construction that translates any network $\mathcal{N}$ with h source messages to a network $\mathcal{N}'$ with $h' > h$ source messages while keeping the parameters of the solution.

Construction B: Let $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, h)$ be a multicast network, with $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and $h \in \mathbb{N}$. For any $h' \in \mathbb{N}$, $h' > h$, we construct the network $\mathcal{N}' = (\mathcal{G}', \sigma', \mathcal{T}', h')$, with $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$, as follows.

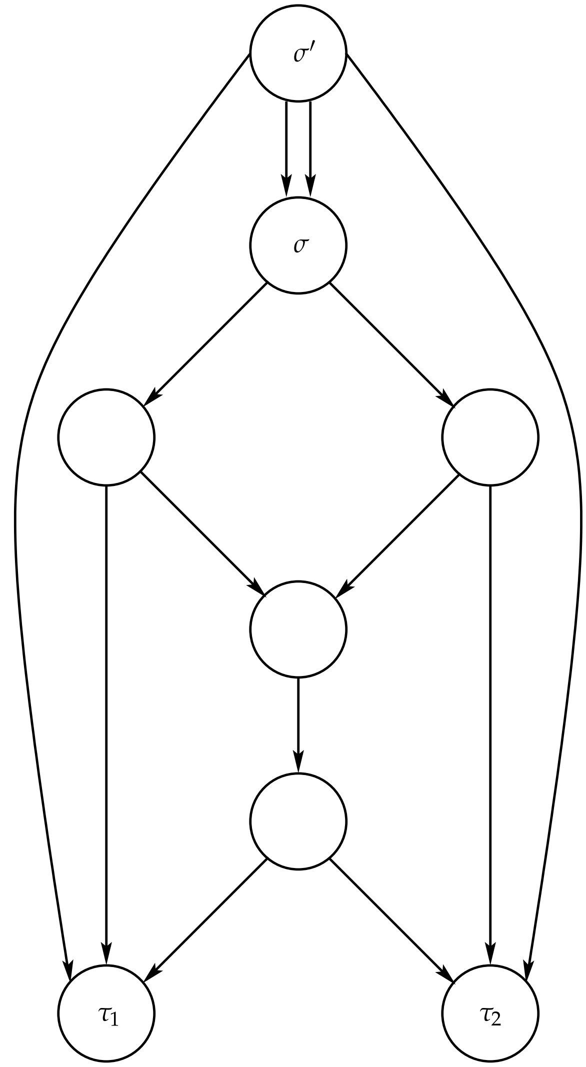

Fig. 4. The result of applying Construction B to the butterfly network of Figure 2, to create a network with the same gap but with three source messages.

We set $\mathcal{V}' = \mathcal{V} \cup \{\sigma'\}$, and $\mathcal{T}' = \mathcal{T}$. We keep all of the edges of $\mathcal{E}$ in $\mathcal{E}'$, and we add h parallel edges from σ' to σ. Finally, for each terminal $\tau \in \mathcal{T}'$, we add $h' - h$ parallel edges from σ' to τ. □

Intuitively, Construction B keeps the network $\mathcal{N}$ in its entirety. The terminals also remain as they were but a new source is added. The new source is connected to the original source by h parallel edges, and to each of the terminals by $h' - h$ parallel edges. We also make the observation that if $\mathcal{N}$ is minimal, then so is $\mathcal{N}'$.

Example 21: We now take the butterfly network, $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, 2)$, from Example 5, and using Construction B, created a minimal network, $\mathcal{N}' = (\mathcal{G}', \sigma', \mathcal{T}, 3)$, with the same gap as $\mathcal{N}$, but with three source messages. The resulting network is depicted in Figure 4.

Lemma 22: Let $h, h' \in \mathbb{N}$, $h' > h$. If $\mathcal{N} = (\mathcal{G}, \sigma, \mathcal{T}, h)$ and $\mathcal{N}' = (\mathcal{G}', \sigma', \mathcal{T}', h')$ are as in Construction B, then there exists a (q,t)-linear solution to $\mathcal{N}$ if and only if there exists a (q,t)-linear solution to $\mathcal{N}'$.

Proof: For the only if part, assume we have a (q,t)-linear solution to $\mathcal{N}$. We build a simple (q,t)-linear solution to $\mathcal{N}'$ by using the same messages over the edges in $\mathcal{E}$, sending the h source messages $x_1,\dots,x_h$ from σ' to σ over the new h edges, and sending $x_{h+1},\dots,x_{h'}$ over the $h'-h$ edges connecting σ' to each of the terminals.

For the if part, assume we have a (q,t)-linear solution to $\mathcal{N}'$. By change of bases we can assume without loss of generality that $x_1,\dots,x_h$ are sent over the h edges connecting σ' to σ. Consider any terminal $\tau\in\mathcal{T}$. Since only $h'-h$ new parallel edges connect σ' to τ directly, the remaining incoming edges into τ, namely, $\mathrm{In}(\tau)\cap\mathcal{E}$, must contain messages that enable the recovery of $x_1,\dots,x_h$. Hence, restricting ourselves to the original network, there is a (q,t)-linear solution allowing σ to convey all of its h messages to the terminals. ■

Corollary 23: For any network $\mathcal{N}=(\mathcal{G},\sigma,\mathcal{T},h)$ there exists a network $\mathcal{N}'=(\mathcal{G}',\sigma',\mathcal{T}',h')$ with $h'>h$ such that

$$\mathrm{gap}(\mathcal{N})=\mathrm{gap}(\mathcal{N}') \quad\text{and}\quad \mathrm{gap}_2(\mathcal{N})=\mathrm{gap}_2(\mathcal{N}').$$

It now follows that there exist networks with $h>2$ source messages that exhibit the same gap as the networks we constructed in the previous section with $h=2$ source messages. Moreover, the minimality of the networks is preserved. However, the networks resulting from Construction B are not combination networks, nor are they sub-networks of combination networks. We therefore return to the main topic of combination networks and their sub-networks, and generalize the previous section's results to more than two source messages.

A full-fledged generalization of the case of $h=2$ to $h>2$ via the skeleton-graph approach, seems highly intricate. Instead, we study Kneser networks $\mathcal{K}_{q,t;h}$ with $h>2$. Their analysis is made possible by replacing their skeleton, a q-analog of the Kneser graph, with a generalization of it to a hypergraph.

We assume throughout this section that $h>2$. We again observe that $\mathcal{K}_{q,t;h}$ is a minimal multicast network which is a sub-network of a minimal combination network of the type $\mathcal{N}_{h,r,h}$. We also observe that trivially,

$$q_v(\mathcal{K}_{q,t;h})\leqslant q^t, \tag{11}$$

since a simple linear solution is for the source σ to send the vector space $V_i\in\begin{bmatrix}V\\t\end{bmatrix}$ to the node in the middle layer indexed by V_i, and all the nodes in the middle layer just forward their incoming message.

It remains to consider a scalar solution for $\mathcal{K}_{q,t;h}$. To that end we define the q-analog Kneser *hypergraph*. The q-analog Kneser hypergraph, denoted $qK^h_{ht:t}$ has $\begin{bmatrix}V\\t\end{bmatrix}=\{V_1,V_2,\dots\}$ as vertices, and an h-hyperedge $\{V_{i_1},V_{i_2},\dots,V_{i_h}\}$ exists, if and only if (7) holds. We note that our generalization to the q-analog Kneser hypergraph is different from other generalizations, e.g., [22] and the references therein.

Several definitions exist for colorings of hypergraphs. In our case, a strong coloring is an assignment of a color to each of the vertices of the hypergraph such that no hyperedge contains two vertices of the same color. The minimal number of colors required to strongly color a given hypergraph $\mathcal{G}$ is called its *strong chromatic number*, and is denoted by $\chi_s(\mathcal{G})$.

Lemma 24: For all $q\in\mathbb{P}$, $t,h\in\mathbb{N}$, and $h\geqslant 3$,

$$\chi_s(qK^h_{ht:t})=\chi(qK_{ht:t}).$$

Proof: We prove every strong coloring of $qK^h_{ht:t}$ is a coloring of $qK_{ht:t}$, and vice versa. First, consider a coloring c of $qK_{ht:t}$. We contend it is also a strong coloring of $qK^h_{ht:t}$. Indeed, for any h subspaces of dimension t, $V_{i_1},\dots,V_{i_h}\subseteq V=\mathbb{F}_q^{ht}$, such that (7) holds, each pair of them is trivially intersecting, hence c gives them all different colors.

In the other direction, assume c is a strong coloring of $qK^h_{ht:t}$ and we prove it is a coloring of $qK_{ht:t}$. Given two trivially intersecting subspaces of dimension t, V_{i_1},V_{i_2}, we can easily build $h-2$ more subspaces of dimension t, $V_{i_3},\dots,V_{i_h}$, such that (7) holds. Thus, V_{i_1} and V_{i_2} are elements in a hyperedge of $qK^h_{ht:t}$, hence c colors them distinctly. ■

We recall more results on the chromatic number of $qK_{n:m}$.

Theorem 25 [5], [7]: Let $q\in\mathbb{P}$, and $n,m\in\mathbb{N}$, $n\geqslant 2m+1$. Except for the case of $n=2m+1$ and $q=2$, in all other cases,

$$\chi(qK_{n:m})=\begin{bmatrix}n-m+1\\1\end{bmatrix}=\frac{q^{n-m+1}-1}{q-1}.$$

We can now state the existence of a gap for Kneser networks with more than two source messages.

Theorem 26: For all $q\in\mathbb{P}$, $t,h\in\mathbb{N}$, $t\geqslant 2$, and $h\geqslant 3$,

$$\begin{aligned}\mathrm{gap}(\mathcal{K}_{q,t;h})&\geqslant\begin{cases}\psi\left(q^t+\frac{1}{h-1}q^{t-1}\right)-q^t & t\geqslant h,\\ \psi\left(q^t+\frac{1}{(h-1)^2}q^{t-1}\right)-q^t & \text{otherwise.}\end{cases}\\ &\geqslant\begin{cases}\frac{1}{h-1}q^{t-1} & t\geqslant h,\\ \frac{1}{(h-1)^2}q^{t-1} & \text{otherwise.}\end{cases}\end{aligned}$$

and

$$\begin{aligned}&\mathrm{gap}_2(\mathcal{K}_{q,t;h})\\ &\quad\geqslant\begin{cases}\log_2\psi\left(q^t+\frac{1}{h-1}q^{t-1}\right)-\log_2 q^t & t\geqslant h,\\ \log_2\psi\left(q^t+\frac{1}{(h-1)^2}q^{t-1}\right)-\log_2 q^t & \text{otherwise.}\end{cases}\\ &\quad\geqslant\begin{cases}\log_2\left(1+\frac{1}{q(h-1)}\right) & t\geqslant h,\\ \log_2\left(1+\frac{1}{q(h-1)^2}\right) & \text{otherwise.}\end{cases}\end{aligned}$$

Proof: As already observed in (11),

$$q_v(\mathcal{K}_{q,t;h})\leqslant q^t.$$

For the scalar solution, our choice of $\mathcal{K}_{q,t;h}$ allows us to follow in a manner similar to the previous section. To avoid tedious notation, denote $q_s\triangleq q_s(\mathcal{K}_{q,t;h})$. The source σ sends a one-dimensional subspace of $\mathbb{F}^h_{q_s}$ on each of the links to the nodes in the middle layer. We may think of the choice of subspace as a color which we assign to the nodes in the middle layer. The total number of colors at our disposal is

$$\begin{bmatrix}h\\1\end{bmatrix}_{q_s}=\frac{q_s^h-1}{q_s-1}.$$

The structure of $\mathcal{K}_{q,t;h}$ implies that a scalar solution must induce a valid strong coloring of $qK^h_{ht:t}$. Therefore, using Lemma 24 and Theorem 25,

$$\frac{q_s^h-1}{q_s-1}\geqslant\chi_s(qK^h_{ht:t})=\chi(qK_{ht:t})=\frac{q^{(h-1)t+1}-1}{q-1}.$$

If we define

$$f(x) \triangleq \frac{x^h - 1}{x - 1} = \sum_{i=0}^{h-1} x^i,$$

then, since $f(x)$ is strictly increasing,

$$q_s \geqslant \sup \left\{ x \in \mathbb{R} \ : \ f(x) \leqslant \frac{q^{(h-1)t+1} - 1}{q - 1} \right\}.$$

If $t \geqslant h$ we observe that for all $1 \leqslant i \leqslant h-1$,

$$\begin{aligned}\left(q^t + \frac{1}{h-1} \cdot q^{t-1}\right)^i &= \sum_{j=0}^{i} \binom{i}{j} \frac{1}{(h-1)^j} \cdot q^{ti-j} \\ &\leqslant \sum_{j=0}^{i} q^{ti-j}\end{aligned}$$

which follows from

$$\binom{i}{j} \frac{1}{(h-1)^j} \leqslant \left(\frac{i}{h-1}\right)^j \leqslant 1,$$

and $i \leqslant h-1$. Thus,

$$\begin{aligned}f\left(q^t + \frac{1}{h-1} \cdot q^{t-1}\right) &= \sum_{i=0}^{h-1} \left(q^t + \frac{1}{h-1} \cdot q^{t-1}\right)^i \\ &\leqslant \sum_{i=0}^{h-1} \sum_{j=0}^{i} q^{ti-j} \\ &\leqslant \sum_{i=0}^{(h-1)t} q^i \\ &= \frac{q^{(h-1)t+1} - 1}{q - 1},\end{aligned}$$

where the last inequality uses the fact that $t \geqslant h$ and hence no power of q repeats in the last double summation. It follows that when $t \geqslant h$ we have

$$q_s \geqslant \psi\left(q^t + \frac{1}{h-1} \cdot q^{t-1}\right).$$

The case of $t < h$ requires a slightly different treatment. For this case, when $1 \leqslant i \leqslant h-1$,

$$\begin{aligned}\left(q^t + \frac{1}{(h-1)^2} \cdot q^{t-1}\right)^i &= \sum_{j=0}^{i} \binom{i}{j} \frac{1}{(h-1)^{2j}} \cdot q^{ti-j} \\ &\overset{(a)}{\leqslant} q^{ti} + \frac{i}{h-1} \cdot q^{ti-1} \\ &\overset{(b)}{\leqslant} q^{ti} + q^{ti-1},\end{aligned}$$

where (a) follows from

$$\binom{i}{j} \frac{1}{(h-1)^{2j}} \leqslant \left(\frac{i}{(h-1)^2}\right)^j \leqslant \frac{1}{h-1},$$

since $i \leqslant h-1$, and (b) follows again from $i \leqslant h-1$. Thus,

$$\begin{aligned}f\left(q^t + \frac{1}{(h-1)^2} \cdot q^{t-1}\right) &= \sum_{i=0}^{h-1} \left(q^t + \frac{1}{(h-1)^2} \cdot q^{t-1}\right)^i \\ &\leqslant 1 + \sum_{i=1}^{h-1} (q^{ti} + q^{ti-1}) \\ &\leqslant \sum_{i=0}^{(h-1)t} q^i \\ &= \frac{q^{(h-1)t+1} - 1}{q - 1}.\end{aligned}$$

It follows that when $t < h$ we have

$$q_s \geqslant \psi\left(q^t + \frac{1}{(h-1)^2} \cdot q^{t-1}\right).$$

By combining the bounds on q_v and q_s in all the cases we obtain the desired result. ■

VI. Minimal Full Combination Networks

We turn to consider, in this section, an upper bound on the gap. In particular, we study the full minimal combination network $\mathcal{N}_{h,r,h}$, and show that it has a very small gap (if at all) compared to its sub-networks, which were considered in the previous section.

The key result we use is the following theorem proved in [25]. Let $(r, q^h, r-s+1)_q$ denote a code over $\mathbb{F}_q$ of length r with q^h codewords and minimum Hamming distance $r-s+1$. If this code is linear, it is denoted by $[r, h, r-s+1]_q$.

Theorem 27 [25]: The $\mathcal{N}_{h,r,s}$ combination network is solvable over $\mathbb{F}_q$ if and only if there exists an $(r, q^h, r-s+1)_q$ code.

In view of Theorem 27, what are the functions on the edges of the $\mathcal{N}_{h,r,s}$ combination network in the three types of solutions – nonlinear, vector linear, and linear?

For the (scalar) *nonlinear* solution, take an $(r, q^h, r-s+1)_q$ code. We can feed the h source messages to an arbitrary encoder for the code to obtain a codeword. The r symbols of the codeword are then transmitted on the r links leaving the source node. In the middle layer, each node simply repeats its incoming message on all of its outgoing links. Finally, each terminal obtains s symbols from the codeword, and since the code has minimum distance of $r-s+1$ it may recover the entire codeword from the surviving s symbols, and reverse the encoding process to obtain the h source messages.

For the scalar *linear* solution, an $[r, h, r-s+1]_q$ code is required. We use the same approach as the one for nonlinear solutions, but using a linear encoder for the code results in a scalar linear solution. Namely, the code has an $r \times h$ generator matrix and the h entries of its ith row are the coding coefficients of the linear function on the link from the source to the ith node in the middle layer.

The vector-linear case differs. A fundamental combinatorial structure that underpins the vector solutions to minimal combination networks is a structure we call a $(t; h, \alpha)_q$-*independent configuration*.

Definition 28: Let $q \in \mathbb{P}$, $t, h, \alpha \in \mathbb{N}$, $\alpha \leqslant h$, and denote $V = \mathbb{F}_q^{ht}$. A $(t; h, \alpha)_q$-*independent configuration (IC)* is a set

$\mathcal{C} = \{V_1, V_2, \dots, V_m\} \subseteq \begin{bmatrix} V \\ t \end{bmatrix}$, such that for all $1 \leqslant i_1 < i_2 < \cdots < i_\alpha \leqslant m$,

$$\dim(V_{i_1} + V_{i_2} + \cdots + V_{i_\alpha}) = \alpha t.$$

We say $|\mathcal{C}| = m$ is the *size* of the IC.

We now make the connection between ICs and full minimal combination networks, $\mathcal{N}_{h,r,h}$.

Lemma 29: The $\mathcal{N}_{h,r,h}$ combination network has a (q,t)-solution if and only if there exists a $(t;h,h)_q$-IC of size r.

Proof: For the only if part, assume that a (q,t)-solution exists. We note that by construction, any node i in the middle layer of $\mathcal{N}_{h,r,h}$ receives a subspace $V_i \subseteq V \triangleq \mathbb{F}_q^{ht}$, with $\dim(V_i) \leqslant t$. If the terminal R_j gets from the middle layer the subspaces $V_{i_1}, \dots, V_{i_h}$, then

$$\dim(V_{i_1} + \cdots + V_{i_h}) = ht,$$

which implies that $\dim(V_i) = t$. Thus, $\{V_i\}_{1 \leqslant i \leqslant r}$ is a $(t;h,h)_q$-IC.

For the if part, assume $\mathcal{C} = \{V_1, \dots, V_r\}$ is a $(t;h,h)_q$-IC. We can easily construct a vector network coding solution to the $\mathcal{N}_{h,r,h}$ combination network. Simply send V_i to the ith middle layer node. Since $\mathcal{C}$ is a $(t;h,h)_q$-IC it follows that each terminal has a full rank $(ht) \times (ht)$ transfer matrix from which it can recover the h messages. ∎

We now bound the size of ICs, which will later enable us to upper bound the gap in minimal combination networks.

Lemma 30: Let $\mathcal{C}$ be a $(t;h,\alpha)_q$-IC. If $\alpha \geqslant 2$ then

$$|\mathcal{C}| \leqslant \frac{q^{(h-\alpha+2)t} - 1}{q^t - 1} + \alpha - 2.$$

Proof: If $\alpha = 2$ the claim is immediate by considering the size of a t-spread in (1).

Assume now $\alpha > 2$, and denote $V \triangleq \mathbb{F}_q^{ht}$. Let us write $\mathcal{C} = \{V_1, V_2, \dots, V_m\}$, and define

$$W_1 \triangleq V_1 + V_2 + \cdots + V_{\alpha-2},$$

where $\dim(W_1) = (\alpha-2)t$. By the definition of an IC, $\mathbb{F}_q^{ht} = W_1 + W_2$, where $W_2 \in \begin{bmatrix} V \\ (h-\alpha+2)t \end{bmatrix}$. It follows that any vector $v \in V_j$, $\alpha - 1 \leqslant j \leqslant m$, may be written uniquely as $v = v_1 + v_2$, where $v_1 \in W_1$ and $v_2 \in W_2$. We now define

$$V_j' \triangleq \{v_2 \ : \ v_1 + v_2 \in V_j,\, v_1 \in W_1,\, v_2 \in W_2\},$$

for all $\alpha - 1 \leqslant j \leqslant m$. It is easily seen that $\dim(V_j') = t$.

Furthermore, for any $\alpha - 1 \leqslant j_1 < j_2 \leqslant m$,

$$\dim(W_1 + V_{j_1}' + V_{j_2}') = \alpha t \Rightarrow \dim(V_{j_1}' + V_{j_2}') = 2t.$$

Thus, the set $\{V_i'\}_{\alpha-1 \leqslant i \leqslant m}$ contains $|\mathcal{C}| - \alpha + 2$ pairwise disjoint t-subspaces of W_2. The number of such subspaces is upper bounded by the size of a t-spread, and thus, by (1),

$$|\mathcal{C}| - \alpha + 2 \leqslant \frac{\begin{bmatrix} (h-\alpha+2)t \\ 1 \end{bmatrix}}{\begin{bmatrix} t \\ 1 \end{bmatrix}}.$$

∎

When $t = 1$, bounding the size of $(1;h,h)_q$-IC is equivalent to finding the longest MDS codes, and hence related to the MDS conjecture. Thus, Lemma 30 forms a generalization for an upper bound on the length of MDS code. The related results for (scalar, $t = 1$) linear codes are given in [19]. Corollary 7 [19, p. 321] asserts that for an $[n,k,n-k+1]_q$ MDS code, we have that $n \leqslant q + k - 1$. This result is strengthened in Theorem 11 [19, p. 326] by using a more complicated proof based on projective geometry. The theorem asserts that if $k \geqslant 3$ and q is odd then $n \leqslant q + k - 2$. Another involved proof for the same result is given for nonlinear codes in [24, pp. 12-13].

We can now give an upper bound on the gap for minimal combination networks.

Theorem 31: For all $h, r \in \mathbb{N}$, $r \geqslant h \geqslant 2$,

$$\begin{aligned} \mathrm{gap}(\mathcal{N}_{h,r,h}) &\leqslant \psi(r-1) - \psi(r-h+1) \leqslant r + h - 3, \\ \mathrm{gap}_2(\mathcal{N}_{h,r,h}) &\leqslant \log_2 \psi(r-1) - \log_2 \psi(r-h+1) \\ &\leqslant 1 + \log_2 \left(1 + \frac{h-2}{r-h+1}\right), \end{aligned}$$

and for all large enough r,

$$\begin{aligned} \mathrm{gap}(\mathcal{N}_{h,r,h}) &\leqslant (r-1)^{21/40} + h - 2, \\ \mathrm{gap}_2(\mathcal{N}_{h,r,h}) &\leqslant \log_2 \left(1 + \frac{(r-1)^{21/40} + h - 2}{r-h+1}\right). \end{aligned}$$

Proof: By [25], a $(q,1)$-scalar linear solution to $\mathcal{N}_{h,r,h}$ exists if and only if an $[r,h,r-h+1]_{q_s}$ MDS code exists. Thus, we certainly have

$$q_s(\mathcal{N}_{h,r,h}) \leqslant \psi(r-1),$$

(e.g., see [19, pp. 317–331]). On the other hand, by Lemma 29, the existence of a (q,t)-solution to $\mathcal{N}_{h,r,h}$ implies the existence of a $(t;h,h)_q$-IC of size r. Then, by Lemma 30,

$$r \leqslant \frac{q^{2t} - 1}{q^t - 1} + h - 2.$$

After rearranging we get,

$$q^t \geqslant r - h + 1,$$

which gives us,

$$q_v(\mathcal{N}_{h,r,h}) \geqslant \psi(r-h+1).$$

Thus, by (3), for all $r \in \mathbb{N}$,

$$\begin{aligned} \mathrm{gap}(\mathcal{N}_{h,r,h}) &\leqslant \psi(r-1) - \psi(r-h+1) \\ &= (\psi(r-1) - (r-1)) \\ &\quad - (\psi(r-h+1) - (r-h+1)) + h - 2 \\ &\leqslant r + h - 3, \end{aligned}$$

and by (4), for all large enough r,

$$\mathrm{gap}(\mathcal{N}_{h,r,h}) \leqslant (r-1)^{21/40} + h - 2.$$

Similarly we get the claimed results for $\mathrm{gap}_2(\mathcal{N}_{h,r,h})$. ∎

Loosely speaking, for the minimal combination network, $\mathcal{N}_{h,r,h}$, the gap is upper bounded by $h-2$ plus the maximum distance between powers of primes. Since it is conjectured that $\psi(n) - n \leqslant O(\log n)$, the upper bound may be further reduced in the future if this conjecture is proved.

VII. Conclusion

In this paper we studied scalar and vector linear solutions to the combination network and its sub-networks. We first found the maximal possible gap for minimal multicast networks with $h = 2$ source messages, and showed that the Kneser network, which was defined in this work, attains this bound on the gap with equality. Furthermore, we studied Kneser networks with $h \geqslant 3$ source messages and proved they also exhibit a gap. Finally, we provided an upper bound on the gap of full minimal combination networks, showing the gap is relatively small.

It is interesting to compare the gap results with those of [9]. In [9] it was proved that for even $h \geqslant 4$ there exists a multicast network $\mathcal{N}$ for which $\mathrm{gap}(\mathcal{N}) = q^{(h-2)t^2/h+o(t)}$, and for odd $h \geqslant 5$ there exists a multicast network $\mathcal{N}$ for which $\mathrm{gap}(\mathcal{N}) = q^{(h-3)t^2/(h-1)+o(t)}$, and in any case, the gap is of order $q^{\Theta(t^2)}$, where q and h are considered constant. In comparison with the results of Section IV and Section V, the gap we present is only of the order of $q^{\Theta(t)}$, which we also prove is optimal in the case of minimal multicast networks with $h = 2$. Two important points distinguish between the two papers: the networks in [9] are not minimal, and [9] uses $h \geqslant 4$. Which of the two, minimality or the number of source messages, contributes to this difference is still an open question. The results of Section V may hint at the former, but we are still missing an upper bound for $h \geqslant 3$.

We also want to highlight the significance of Section VI. If we take for example the case of $h = 2$ source messages, by Theorem 31, the full minimal combination network $\mathcal{N}_{2,r,2}$ has 0 gap, for any $r \in \mathbb{N}$ (a fact that was noticed in [9] using Latin squares). If we pick $q \in \mathbb{P}$, $t \in \mathbb{N}$, $r = \begin{bmatrix} 2t \\ t \end{bmatrix}$, keep the source and all of the middle layer of nodes, but remove a few of the terminals, we obtain the Kneser network $\mathcal{K}_{q,t;2}$, the network from Theorem 15, that has the maximal possible gap, $\psi(q^t + q^{t-1} - 1) - q^t$. Thus, it appears that $\mathcal{N}_{2,r,2}$ is equally difficult for a scalar solution as it is for a vector solution. However, by pruning some terminals, the resulting Kneser network $\mathcal{K}_{q,t;2}$ becomes easy for a vector solution but difficult for a scalar solution. This upper bound on the gap by Theorem 31 grows very slowly. As another example, if $r-2$ is not a prime power then $\mathcal{N}_{3,r,3}$ has no gap. As a consequence, from an engineering point of view, it seems beneficial to use the sub-networks of the combination network, whose skeleton is a q-Kneser graph. Thus, on the one hand, the network has a simple three-layer structure as the combination networks, but on the other hand, provides a maximum efficiency in terms of field size, when using vector-linear network coding.

A full characterization of the gap in linear multicast networks is still an open question. We also suggest minimality of such networks may play a role in limiting the gap, a problem which we leave for future work. Finally, the gap between non-linear and linear solutions for multicast networks is still unknown.

Acknowledgment

The authors would like to thank Y. Zhang for helpful discussions.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[2] T. M. Apostol, *Introduction to Analytic Number Theory*. New York, NY, USA: Springer-Verlag, 1976.

[3] R. C. Baker, G. Harman, and J. Pintz, "The difference between consecutive primes, II," *Proc. London Math. Soc.*, vol. 83, no. 3, pp. 532–562, Nov. 2001.

[4] M. Blaum, J. Bruck, and A. Vardy, "MDS array codes with independent parity symbols," *IEEE Trans. Inf. Theory*, vol. 42, no. 2, pp. 529–542, Mar. 1996.

[5] A. Blokhuis *et al.*, "A Hilton-Milner theorem for vector spaces," *Electron. J. Combinatorics*, vol. 17, no. 1, pp. 1–12, May 2010.

[6] A. Blokhuis, A. E. Brouwer, and T. Szőnyi, "On the chromatic number of q-Kneser graphs," *Des., Codes Cryptogr.*, vol. 65, no. 3, pp. 187–197, 2012.

[7] A. Chowdhury, C. Godsil, and G. Royle, "Colouring lines in projective space," *J. Combinat. Theory A*, vol. 113, no. 1, pp. 39–52, Jan. 2006.

[8] S. Y. el Rouayheb, C. N. Georghiades, and A. Sprintson, "Network coding in minimal multicast networks," in *Proc. IEEE Inf. Theory Workshop*, Punta del Este, Uruguay, Mar. 2006, pp. 307–311.

[9] T. Etzion and A. Wachter-Zeh, "Vector network coding based on subspace codes outperforms scalar linear network coding," *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 2460–2473, Apr. 2018.

[10] C. Fragouli and E. Soljanin, "Information flow decomposition for network coding," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 829–848, Mar. 2006.

[11] C. Fragouli and E. Soljanin, "(Secure) linear network coding multicast: A theoretical minimum and some open problems," *Des., Codes Cryptography*, vol. 78, no. 1, pp. 269–310, Jan. 2016.

[12] P. Frankl and R. M. Wilson, "The Erdös-Ko-Rado theorem for vector spaces," *J. Combin. Theory, A*, vol. 43, pp. 228–236, Nov. 1986.

[13] S. Gheorghiu, S. S. Bidokhti, C. Fragouli, and A. L. Toledo, "Degraded multicasting with network coding over the combination network," in *Proc. Int. Symp. Netw. Coding*, Beijing, China, Jul. 2011, pp. 1–6.

[14] W. Guo, X. Shi, N. Cai, and M. Medard, "Localized dimension growth: A convolutional random network coding approach to managing memory and decoding delay," *IEEE Trans. Commun.*, vol. 61, no. 9, pp. 3894–3905, Sep. 2013.

[15] F. Ihringer, "The chromatic number of the q-Kneser graph for large q," *Electron. J. Combinatorics*, vol. 26, no. 1, pp. 1–12, Mar. 2019.

[16] S. Jaggi *et al.*, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1973–1982, Jun. 2005.

[17] A. R. Lehman and E. Lehman, "Complexity classification of network information flow problems," in *Proc. 15th Annu. ACM-SIAM Symp. Discrete Algorithms*, New Orleans, LA, USA, Jan. 2004, pp. 142–150.

[18] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inf. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.

[19] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error Correcting Codes*. Amsterdam, The Netherlands: North Holland, 1978.

[20] S. Maheshwar, Z. Li, and B. Li, "Bounding the coding advantage of combination network coding in undirected networks," *IEEE Trans. Inf. Theory*, vol. 58, no. 2, pp. 570–584, Feb. 2012.

[21] M. Médard, M. Effros, T. Ho, and D. Karger, "On coding for non-multicast networks," in *Proc. 41st Allerton Conf. Commun., Control, Comput.*, Monticello, IL, USA, Oct. 2003, pp. 21–29.

[22] F. Meunier, "Colorful subhypergraphs in kneser hypergraphs," *Electron. J. Combinatorics*, vol. 21, no. 1, pp. 1–13, Jan. 2014.

[23] C. K. Ngai and R. W. Yeung, "Network coding gain of combination networks," in *Proc. IEEE Inf. Theory Workshop (ITW)*, San Antonio, TX, USA, 2004, pp. 283–287.

[24] D. Raghavarao, *Constructions and Combinatorial Problems in Design of Experiments*. Hoboken, NJ, USA: Wiley, 1971.

[25] S. Riis and R. Ahlswede, "Problems in network coding and error correcting codes," in *General Theory of Information Transfer and Combinatorics* (Lecture Notes in Computer Science), vol. 4123. Berlin, Germany: Springer, 2006, pp. 861–897.

[26] M. Schwartz and T. Etzion, "Codes and anticodes in the grassman graph," *J. Combinat. Theory A*, vol. 97, no. 1, pp. 27–42, Jan. 2002.

[27] Q. T. Sun, X. Yang, K. Long, X. Yin, and Z. Li, "On vector linear solvability of multicast networks," *IEEE Trans. Commun.*, vol. 64, no. 12, pp. 5096–5107, Dec. 2016.

[28] M. Xiao, M. Medard, and T. Aulin, "A binary coding approach for combination networks and general erasure networks," in *Proc. IEEE Int. Symp. Inf. Theory*, Nice, France, Jun. 2007, pp. 786–790.
[29] R. W. Yeung, S.-Y. R. Li, N. Cai, and Z. Zhang, *Network Coding Theory*. Boston, MA, USA: Now, 2005.

Han Cai (Member, IEEE) received the B.S. and M.S. degrees in mathematics from Hubei University, Wuhan, China, in 2009 and 2013, respectively, and the Ph.D. degree from the Department of Communication Engineering, Southwest Jiaotong University, Chengdu, China, in 2017. From October 2015 to October 2017, he was a Visiting Ph.D. Student with the Faculty of Engineering, Information and Systems, University of Tsukuba, Japan. He is currently a Post-Doctoral Fellow with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Israel. His research interests include coding theory and sequence design.

Johan Chrisnata received the bachelor's degree in mathematics from Nanyang Technological University (NTU), Singapore, in 2015. He is currently pursuing the joint Ph.D. degree in mathematics with the School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore, and in computer science with the Department of Computer Science, Technion, Israel. From August 2015 until August 2018, he was a Research Officer with NTU. His research interests include enumerative combinatorics and coding theory.

Tuvi Etzion (Fellow, IEEE) was born in Tel Aviv, Israel, in 1956. He received the B.A., M.Sc., and D.Sc. degrees from the Technion–Israel Institute of Technology, Haifa, Israel, in 1980, 1982, and 1984, respectively.

In 1984, he held a position at the Department of Computer Science, Technion, where he is currently the Bernard Elkin Chair of computer science. From 1985 to 1987, he was a Visiting Research Professor with the Department of Electrical Engineering—Systems, University of Southern California, Los Angeles. In the summers of 1990 and 1991, he was a visiting Bellcore, Morristown, NJ, USA. From 1994 to 1996, he was a Visiting Research Fellow with the Computer Science Department, Royal Holloway University of London, Egham, U.K. He also had several visits to the Coordinated Science Laboratory, University of Illinois in Urbana-Champaign, from 1995 to 1998, two visits to HP Bristol in the summers of 1996 and 2000, a few visits to the Department of Electrical Engineering, University of California at San Diego, from 2000 to 2017, several visits to the Mathematics Department, Royal Holloway University of London, from 2007 to 2017, a few visits to the School of Physical and Mathematical Science (SPMS), Nanyang Technological University, and to the Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore, from 2016 to 2019, and a few visits to Jiaotong University, Beijing, from 2017 to 2019. His research interests include applications of discrete mathematics to problems in computer science and information theory, coding theory, network coding, and combinatorial designs.

Dr. Etzion was an Associate Editor for Coding Theory of the IEEE TRANSACTIONS ON INFORMATION THEORY from 2006 till 2009. From 2004 to 2009, he was an Editor of the *Journal of Combinatorial Designs*. Since 2011, he has been an Editor of *Designs, Codes, and Cryptography*. Since 2013, he has also been an Editor of *Advances in Mathematics of Communications*.

Moshe Schwartz (Senior Member, IEEE) received the B.A. *(summa cum laude)*, M.Sc., and Ph.D. degrees from the Computer Science Department, Technion—Israel Institute of Technology, Haifa, Israel, in 1997, 1998, and 2004, respectively.

He was a Fulbright Post-Doctoral Researcher with the Department of Electrical and Computer Engineering, University of California San Diego, and a Post-Doctoral Researcher with the Department of Electrical Engineering, California Institute of Technology. While on sabbatical 2012–2014, he was a Visiting Scientist with the Massachusetts Institute of Technology (MIT). He is currently a Professor with the School of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Israel. His research interests include algebraic coding, combinatorial structures, and digital sequences.

Prof. Schwartz received the 2009 IEEE Communications Society Best Paper Award in Signal Processing and Coding for Data Storage, and the 2020 NVMW Persistant Impact Prize. He has been serving as an Associate Editor for Coding Techniques for the IEEE TRANSACTIONS ON INFORMATION THEORY since 2014.

Antonia Wachter-Zeh (Member, IEEE) received the M.Sc. degree in communications technology from Ulm University, Germany, in 2009, and the Ph.D. degree from the Institute of Communications Engineering, Ulm University, Germany, and from the Institut de recherche mathematique de Rennes (IRMAR), Universite de Rennes 1, Rennes, France, in 2013. From 2013 to 2016, she was a Post-Doctoral Researcher with the Technion–Israel Institute of Technology, Haifa, Israel. She is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Technical University of Munich (TUM), Munich Germany. Her research interests are coding theory, information theory, and cryptography and their application to storage, communications, and security. She was a recipient of the DFG Heinz Maier-Leibnitz-Preis and of an ERC Starting Grant.