

Quasi-linear Network Coding

Moshe Schwartz

Electrical and Computer Engineering
Ben-Gurion University of the Negev
Beer Sheva 8410501, Israel
schwartz@ee.bgu.ac.il

Muriel Médard

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
medard@mit.edu

Abstract—We present a heuristic for designing vector non-linear network codes for non-multicast networks, which we call quasi-linear network codes. The method presented has two phases: finding an approximate linear network code over the reals, and then quantizing it to a vector non-linear network code using a fixed-point representation. Apart from describing the method, we draw some links between some network parameters and the rate of the resulting code.

I. INTRODUCTION

Network coding was introduced in [1] as a means of increasing the amount of information flowing through a network. In this scheme, a network is a directed graph, where information is generated by source nodes, and demanded by terminal nodes. All participating nodes receive information through their incoming edges, combine the information, and transmit it over their outgoing edges.

Linear network coding has drawn particular interest due its simplicity and structure. Works such as [1], [5], [6], [8] studied fundamental bounds on the parameters of such codes, mainly for linear multicast networks. For these networks, necessary and sufficient conditions for the existence of a linear solution are known, as are efficient algorithms for finding such a solution. It was also shown in [9], that solvable multicast networks always have a scalar linear solution.

In the non-multicast case, the picture is more complicated. Several works [10], [12], [13] showed various restrictions on the ability to find a linear solution to these networks. These culminated in [3], that showed linear network codes are insufficient for solving a non-multicast network over finite fields, commutative rings, and even R -modules. This was done using a single network, for which no linear solution exists, though an ad-hoc non-linear solution is possible. Certain classes of networks are known to have linear solutions over finite fields for the non-multicast case. The necessity, for the existence of a scalar linear solution over a finite field, of the matroidal nature of a network was shown in [4], and the sufficiency proven in [7]. Recently, [11] established the equivalence between the discrete polymatroidal nature of a network and the existence of a vector linear solution over a finite field for all feasible non-multicast connections over that network.

The goal of this paper is to introduce *quasi-linear network coding*, which is a heuristic method for designing a vector non-linear network code for non-multicast networks. The method is inspired by work on real network coding of [14]. The method we suggest has two main phases. In the first one, an

approximate solution for the network is found over the reals. We say this is an *approximate* solution, since at the terminals, the original messages are not recovered exactly, and there is some mixing with unwanted messages. At the second stage, restrictions over possible source messages, together with a fixed-point representation, enable the terminals to reconstruct the demanded source messages with zero error. We thus gain from both worlds: the method is linear at its core, giving it some structure, while at the second phase non-linearity is introduced in a systematic way, overcoming the insufficiency of linear solutions to non-multicast networks.

Comparing this work with [14], we note that both works try to solve the network over the reals. However, [14] considers only the multicast case, and it assumes an exact solution exists. Furthermore, the real coefficients in [14] are used to obtain a graceful degradation. This is inherently different in the non-multicast case that we consider, since an exact solution is not guaranteed, in which case we use the real coefficients to obtain an approximate solution.

This paper is organized as follows. In Section II we introduce the required definition and notation. In Section III we describe quasi-linear network codes, discuss some of their properties, and give examples. We conclude in Section IV with a brief summary and some open questions.

II. PRELIMINARIES

For the purpose of this work, a *network* is a directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. For a vertex $v \in \mathcal{V}$ we denote its incoming edges as $\text{In}(v)$, and its outgoing edges as $\text{Out}(v)$. The in-degree of vertex $v \in \mathcal{V}$ is defined as $\delta(v) = |\text{In}(v)|$, and the maximal in-degree in the graph is denoted as

$$\delta = \max_{v \in \mathcal{V}} \delta(v).$$

The *depth* of the graph, denoted d , is defined as the length of the longest path in the graph. Since the graph is acyclic, the depth is well defined.

Two distinguished subsets of vertices are the *source nodes* and the *terminal nodes*, denoted by $S, T \subseteq \mathcal{V}$ respectively. The source nodes generate *messages*, which are symbols from some finite alphabet Σ . With each of the terminal nodes we associate a demand for a subset of the messages.

Information is transmitted over edges in the form of symbols from Σ . We denote this transmitted information as $\text{val}(e) \in \Sigma$ for each $e \in \mathcal{E}$. Apart from the source nodes, each node $v \in \mathcal{V}$, transmits information along its outgoing edges, which is a function of the information received on the incoming

edges to v . More precisely, for each $v \in \mathcal{V} \setminus S$, and for each $e \in \text{Out}(v)$,

$$\text{val}(e) = f_e(e_1, e_2, \dots, e_{\delta(v)}),$$

where $\text{In}(v) = \{e_1, \dots, e_{\delta(v)}\}$, and f_e is a function associated with the edge e . Here we implicitly assume a fixed order of the edges in $\text{In}(v)$, since f_e is not necessarily a symmetric function.

Throughout this paper we assume a single source node, i.e., $S = \{s\}$. The terminal nodes are denoted $T = \{t_1, t_2, \dots, t_\ell\}$. To avoid trivialities, since the graph is acyclic and there is only one source node, we can assume s is the only vertex in \mathcal{G} with no incoming edges. Thus, the outgoing edges of s transmit the messages of s .

The case we study is a general *non-multicast network*, where each of the terminal nodes demands a subset of the source messages. The subsets are not necessarily disjoint.

We say a vertex $v \in \mathcal{V}$ is of depth $d(v)$ if the *longest* path from s to v is of length d . In a similar manner, the depth of an edge $e \in \mathcal{E}$, $e = v \rightarrow v'$, is defined as the depth of v , and we denote $d(e) = d(v)$. We can now partition the edge set

$$\mathcal{E} = \mathcal{E}_0 \cup \mathcal{E}_1 \cup \dots \cup \mathcal{E}_{d-1}.$$

By definition, $e \in \mathcal{E}_{d(e)}$ for each $e \in \mathcal{E}$. We note that $\mathcal{E}_0 = \text{Out}(s)$. Again, in order to avoid trivialities, we assume \mathcal{E}_{d-1} contains at least one edge ending in a terminal node.

III. METHOD AND ANALYSIS

The quasi-linear network-coding method we describe is inspired by the arithmetic network coding of [14]. While the latter work considered multicast setting, we consider the general non-multicast which does subsume the multicast case. The main strategy is given by the following two steps:

- 1) Initially, instead of using a finite field for the alphabet of messages, we use real numbers. Nodes linearly combine the real scalars on incoming edges using real coefficients. The coefficients are chosen so as to approximate the demands at the terminal nodes.
- 2) Messages from the source are restricted to integers. The messages over the edges are replaced with finite-precision fixed-point representations in base b . The degree of approximation to the demands, calculated in the first step, is used to limit the range of integers the source may send. The terminal nodes reverse the linear combination and quantize the result to the nearest integer, as an estimate to the demanded messages.

We observe that at the terminal nodes there are two sources of noise that may prevent the recovery of the correct demanded messages. The first is due to the numerical error accumulating along the paths from the source to the terminal, caused by the restriction to fixed-point precision. This occurs even if we assume internal computations within the node are done with infinite precision¹.

¹There is no actual need for infinite precision. We can choose a precision high enough within nodes to make internal node computations irrelevant.

The second source of noise at the terminals is due to the approximation to the demand. Terminal nodes essentially compute a linear combination of the messages from the source. Ideally, this combination has a coefficient of 1 for the demanded message, and a coefficient of 0 for each of the other messages. However, such a solution may not be possible, as was demonstrated in [3]. We shall therefore strive to obtain a linear combination with coefficients close to 1 and 0 appropriately. Intuitively, such combinations introduce a “weak” version of unwanted messages. By limiting the range of messages the source transmits, these “weak” versions of unwanted messages lead to interference that is removed by quantization. Thus, when the terminal recovers the correct integer message when quantizing the linear combination to the nearest integer.

A. Stage I – Working over \mathbb{R}

We now describe the method in detail, and analyze some of its properties. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a non-multicast network, with source node s and terminal nodes $T = \{t_1, \dots, t_\ell\}$. Let $k = |\text{Out}(s)|$ be the number of outgoing edges from the source s , and let $m_1, m_2, \dots, m_k \in \mathbb{Z}$ be the messages the source transmits, each on one of its outgoing edges, $\text{Out}(s) = \mathcal{E}_0$. We conveniently denote $\text{Out}(s) = \{e_1^s, e_2^s, \dots, e_k^s\}$, and thus, $\text{val}(e_i^s) = m_i$ for all $i \in [k]$.

For the first stage of the design of the algorithm, assume that over the rest of the edges, $\mathcal{E} \setminus \mathcal{E}_0$, values from \mathbb{R} are transmitted. We denote the value transmitted over edge $e \in \mathcal{E}$ as $\text{val}(e)$. Every vertex in the graph transmits, over its outgoing edges, linear combinations of values received from the incoming edges. To be more precise, for every $v \in \mathcal{V}$, and for every $e' \in \text{Out}(v)$, the value transmitted over e' is the linear combination

$$\text{val}(e') = \sum_{e \in \text{In}(v)} \alpha_{e \rightarrow e'} \text{val}(e),$$

for some fixed coefficients $\alpha_{e \rightarrow e'} \in \mathbb{R}$. For our convenience, if $e \rightarrow e'$ is not a path in \mathcal{G} , we shall set $\alpha_{e \rightarrow e'} = 0$.

As is usually done in network coding, let T be the $|\mathcal{E}| \times |\mathcal{E}|$ matrix of the line-graph of \mathcal{G} , with

$$T_{e, e'} = \alpha_{e \rightarrow e'}.$$

It is well-known [8] that by running the network-coding system, the value transmitted over any edge $e \in \mathcal{E}$ is given by

$$\text{val}(e) = \sum_{i=1}^k m_i \left(\sum_{j=0}^d T^j \right)_{e_i^s, e}$$

where $T^0 = I$, the identity matrix.

For the sake of brevity and ease of notation, let us assume each of the terminals demands a single message from the source. The more general case is a trivial extension of the case we describe. Say terminal $t \in T$ demands the w_t th source message, i.e., m_{w_t} , where $w_t \in [k]$. To that end, the terminal t chooses $\delta(t)$ real coefficients, $\beta_{t,1}, \dots, \beta_{t,\delta(t)}$, each associated

with the $\delta(t)$ incoming edges, denoted $e'_1, \dots, e'_{\delta(t)}$. The terminal node t then performs the linear combination

$$\sum_{r=1}^{\delta(t)} \beta_{t,r} \text{val}(e'_r) = \sum_{i=1}^k m_i \sum_{r=1}^{\delta(t)} \left(\sum_{j=0}^d T^j \right)_{e'_i, e'_r}. \quad (1)$$

We denote

$$\gamma_{t,i} = \sum_{r=1}^{\delta(t)} \left(\sum_{j=0}^d T^j \right)_{e'_i, e'_r}, \quad (2)$$

and then we can rewrite (1) as

$$\sum_{r=1}^{\delta(t)} \beta_{t,r} \text{val}(e'_r) = \sum_{i=1}^k \gamma_{t,i} m_i. \quad (3)$$

Since terminal t demands the w_t th message, ideally we would like to get $\gamma_{t,w_t} = 1$ and $\gamma_{t,i} = 0$ for all $i \neq w_t$. This goal may be unattainable, in particular, since we need to solve this concurrently for all $t \in T$.

We therefore resort to try and find an approximate solution over the reals, as follows. We define the function

$$\mathcal{F} = \sum_{t \in T} \left((\gamma_{t,w_t} - 1)^2 + \sum_{\substack{i \in [\delta(t)] \\ i \neq w_t}} \gamma_{t,i}^2 \right),$$

where $\gamma_{t,i}$ is defined in (2). By choosing the real coefficients $\alpha_{e \rightarrow e'}$ and $\beta_{t,i}$, the goal is to minimize \mathcal{F} .

As a crude overall measure of approximation, we define γ as

$$\gamma = \max_{t \in T} \left(|\gamma_{t,w_t} - 1| + \sum_{\substack{i \in [\delta(t)] \\ i \neq w_t}} |\gamma_{t,i}| \right).$$

Intuitively, γ is the maximal magnitude of deviation from the coefficients being 1 or 0 appropriately. We say the real solution to the network is *exact* if $\gamma = 0$. This concludes the first stage of designing a quasi-linear network code.

B. Stage II – Working with Fixed-Point Precision

The goal of the second stage of the design of quasi-linear network codes, is to quantize all the real numbers transmitted over edges to a fixed-point presentation in base b . By doing so, we enable the transmission of real values as symbols from a finite alphabet, but also introduce more noise into the system. In this section we go through this quantization process, and bound the amount of noise introduced. This will be helpful in determining the range of possible messages that can be recovered with zero error at the terminals. It should be noted that the bounds we give are general but crude, and that for specific networks, a careful analysis in the spirit of this section, will provide better bounds.

We denote the maximal linear-combination coefficient, chosen in the previous stage, as

$$\alpha = \max_{e, e' \in \mathcal{E}} |\alpha_{e \rightarrow e'}|.$$

We assume $\alpha > 1$. If that is not the case, either the solution is a trivial routing, or we can scale all the coefficients and receive a scaled version of the result at the terminals. We also assume that the magnitude of all source messages is upper bounded by

$$|m_i| \leq M.$$

Using our partition of the edges by depth, let us denote

$$M_i = \max_{e \in \mathcal{E}_i} |\text{val}(e)|.$$

Lemma 1. For all $0 \leq i \leq d-1$ we have

$$M_i \leq (\delta\alpha)^i M.$$

Proof: This is a simple proof by induction. For the induction base we have for each $e \in \mathcal{E}_0$

$$|\text{val}(e)| \leq M = (\delta\alpha)^0 M.$$

Since this is true for all edges $e \in \mathcal{E}_0$ we have

$$M_0 \leq (\delta\alpha)^0 M.$$

For the induction step, let $e' = v \rightarrow v'$ be some edge, $e' \in \mathcal{E}_i$. Then

$$\begin{aligned} |\text{val}(e')| &= \left| \sum_{e \in \text{In}(v)} \alpha_{e \rightarrow e'} \text{val}(e) \right| \leq \sum_{e \in \text{In}(v)} |\alpha_{e \rightarrow e'} \text{val}(e)| \\ &\leq \sum_{e \in \text{In}(v)} \alpha M_{d(e)} \leq \sum_{e \in \text{In}(v)} \alpha (\delta\alpha)^{d(e)} M \\ &\leq \sum_{e \in \text{In}(v)} \alpha (\delta\alpha)^{d(e)-1} M = (\delta\alpha)^i M, \end{aligned}$$

where we used the fact that $\alpha\delta \geq 1$. Since this holds for any $e \in \mathcal{E}_i$ we have

$$M_i \leq (\delta\alpha)^i M. \quad \blacksquare$$

Suppose now every edge can carry a value in fixed-point base- b representation with P digits left of the fixed point, and p digits to the right of it. The nodes calculate the same linear combinations on their inputs as before, and we assume infinite precision within the nodes. However, before transmitting the results over the outgoing messages, a quantization occurs.

We denote the value of the fixed-point representation sent over edges as $\widetilde{\text{val}}(e)$. For all $0 \leq i \leq d-1$ we denote

$$\epsilon_i = \max_{e \in \mathcal{E}_i} \left| \widetilde{\text{val}}(e) - \text{val}(e) \right|.$$

Lemma 2. For all $0 \leq i \leq d-1$ we have

$$\epsilon_i \leq \frac{(\delta\alpha)^i - 1}{\delta\alpha - 1} b^{-p}.$$

Proof: For convenience, let us denote the RHS of the claim as $f(i)$. We observe that $f(i)$ is non-negative and monotone increasing in i .

The proof is by induction. For the induction base we note that $\epsilon_0 = 0$ since, in our setting, the source node transmits

only integer values, and we shall make sure to set P to a large enough value so that no truncation error occurs.

For the induction step consider any edge $e' \in \mathcal{E}_i$. We now have

$$\begin{aligned}
 & \left| \widetilde{\text{val}}(e') - \text{val}(e') \right| = \\
 & = b^{-P} + \left| \sum_{e \in \text{In}(v)} \alpha_{e \rightarrow e'} \widetilde{\text{val}}(e) - \alpha_{e \rightarrow e'} \text{val}(e) \right| \\
 & \leq b^{-P} + \sum_{e \in \text{In}(v)} \left| \alpha_{e \rightarrow e'} \epsilon_d(e) \right| \\
 & \leq b^{-P} + \sum_{e \in \text{In}(v)} |\alpha_{e \rightarrow e'} f(d(e))| \\
 & \leq b^{-P} + \sum_{e \in \text{In}(v)} \alpha f(i-1) \\
 & \leq b^{-P} + \delta \alpha f(i-1) \\
 & = f(i).
 \end{aligned}$$

Since this holds for any $e' \in \mathcal{E}_i$ we have

$$\epsilon_i \leq f(i).$$

We are now in a position to combine all of the previous observations, and give sufficient conditions for a zero-error recovery of the values at the terminals.

Theorem 3. *Using the quasi-linear network-coding scheme describe before, it is possible to recover the original values at the terminals if*

$$\begin{aligned}
 M &< \frac{1}{2\gamma} \quad (\text{only if } \gamma > 0), \\
 p &> \log_b \frac{(\delta\alpha)^{d-1} - 1}{\delta\alpha - 1} - \log_b \left(\frac{1}{2} - \gamma M \right), \\
 P &\geq \log_b \left(2(\delta\alpha)^{d-1} M + 2 \right).
 \end{aligned}$$

Proof: We first note that at any terminal $t \in T$, the absolute difference between the demanded message and the linear combination obtained by (3) and Lemma 2, is upper bounded by

$$\left| m_{w_t} - \sum_{i=1}^k \gamma_{t,i} m_i \right| + \epsilon_{d-1} \leq \gamma M + \frac{(\delta\alpha)^{d-1} - 1}{\delta\alpha - 1} b^{-P}.$$

Thus, if we require

$$\gamma M + \frac{(\delta\alpha)^{d-1} - 1}{\delta\alpha - 1} b^{-P} < \frac{1}{2},$$

then rounding the resulting linear combination at terminal t to the nearest integer, will recover the message correctly.

It follows that when $\gamma > 0$, i.e., the real solution in the first stage is not exact, we must require

$$\gamma M < \frac{1}{2}.$$

Furthermore, after rearranging and solving for p , we obtain the desired requirement,

$$p > \log_b \frac{(\delta\alpha)^{d-1} - 1}{\delta\alpha - 1} - \log_b \left(\frac{1}{2} - \gamma M \right).$$

Furthermore, according to Lemma 1, the maximal value sent on an edge is upper bounded by

$$M_{d-1} = (\delta\alpha)^{d-1} M.$$

Since, in the previous paragraphs, we bounded the quantization error on any edge by $\frac{1}{2}$, and since we need the integers in the range $[-M_{d-1} - \frac{1}{2}, M_{d-1} + \frac{1}{2}]$, taking $P \geq \log_b \left(2(\delta\alpha)^{d-1} M + 2 \right)$ digits to the left of the fixed point ensures the value is within the representation range. ■

We briefly pause to contemplate the implications of Theorem 3. If, in the first stage, we are able to find an *exact* real solution, i.e., with $\gamma = 0$, then there is no bound on the magnitude of the source messages sent. It follows that, in this case, the number of digits used when transmitting over any edge, $P + p$, is $\log_b M + c$, where c is some constant that depends on the network topology. Thus, we can get quasi-linear network-coding solutions with rate arbitrarily close to 1, where rate is measured as the ration between the minimum number of bits required to describe a source message, and the number of bits used for transmission over any edge.

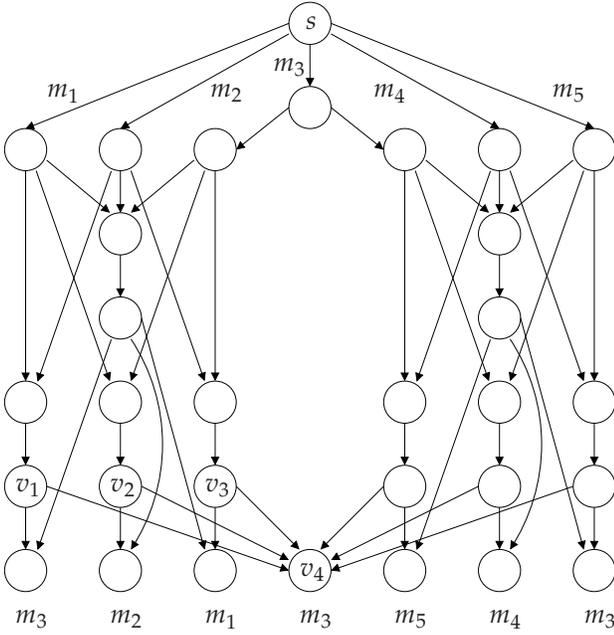
When we do not have an exact real solution in the first stage, we can no longer support arbitrarily-large source messages. Furthermore, we note that the bound on p from Theorem 3 distinctly exhibits a component affected by the degree of approximation γ , and a component affected by the fixed-point quantization.

Example 4. *The network \mathcal{G}_1 , shown in Figure 1, was given in [3] as part of a larger network. The network has a source node s that produces five source messages m_1, \dots, m_5 . There are seven terminal nodes, with single-message demands written beneath the appropriate node.*

It was shown in [3] that \mathcal{G}_1 has no scalar linear network-coding solution over $GF(2^h)$, for any h . If we restrict ourselves to a routing solution, in which nodes cannot linearly combine incoming messages, the best rate we can achieve is $\frac{1}{3}$, as was shown in [2].

However, the network is exactly solvable over \mathbb{R} in the following simple manner. All source nodes repeat their message on all outgoing edges. All internal nodes sum all incoming messages and transmit the sum over all the outgoing edges. The terminal nodes perform simple subtraction to obtain their demands, except for the middle terminal node, v_4 , which computes $\frac{1}{2}(-v_1 + v_2 + v_3)$. It easily follows that given source messages of n bits, we can find a quasi-linear network-coding solution with $P = n + 2$ and $p = 0$ digits to the left and to the right of the fixed point, respectively. Since n is arbitrarily large, the achievable rate is $\frac{n}{n+2}$, which is asymptotically 1 as $n \rightarrow \infty$. □

Example 5. *A more interesting example is the network \mathcal{G}_2 , shown in Figure 2, which was also given in [3]. The network*


 Figure 1. The network \mathcal{G}_1 .

has a single source s , which produces three message m_1 , m_2 , and m_3 . The network also has three terminal nodes, whose demands are written below them.

It was shown in [3], that \mathcal{G}_2 has no scalar linear solution over $GF(q)$ when q is odd, and does have a scalar linear solution over $GF(2^h)$.

We bring this network as an example for a network that has no exact real solution. An approximate solution which was found using a computer search is detailed below:

$$\begin{aligned} \alpha_{e_1 \rightarrow e_5} &= 0.0332528 & \alpha_{e_2 \rightarrow e_5} &= -11.8712 \\ \alpha_{e_3 \rightarrow e_6} &= 16.3384 & \alpha_{e_4 \rightarrow e_6} &= 2.69746 \\ \alpha_{e_7 \rightarrow e_{11}} &= 2.79007 & \alpha_{e_8 \rightarrow e_{11}} &= 2.02721 \\ \alpha_{e_9 \rightarrow e_{12}} &= -1.16509 & \alpha_{e_{10} \rightarrow e_{12}} &= 2.28349 \\ \beta_1 &= -0.0169705 & \beta_2 &= 0.182872 \\ \beta_3 &= -0.030174 & \beta_4 &= 0.0722992 \\ \beta_5 &= -25.8106 & \beta_6 &= 21.8495 \end{aligned}$$

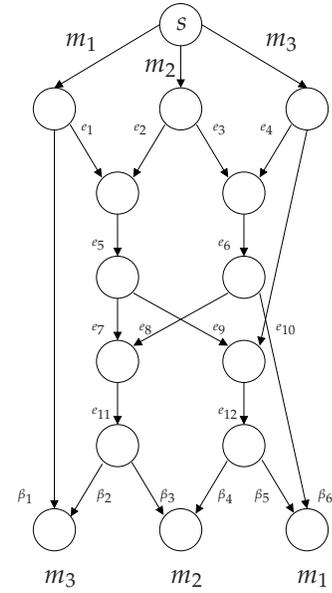
where the β 's are the coefficients used at the terminals to recover the original messages, and are written next to the edge they apply to. All the nodes with a single incoming edge simply repeat the incoming message on all outgoing edges.

The approximation factor turn out to be

$$\gamma = 0.00572545,$$

which, by Theorem 3, allows us to use integers in the range $[-87, 87]$. Since we use base 2 in this example, we use the range $[-64, 63]$, whose integers may be expressed using 7 bits.

The other parameters involved in this network are maximum in-degree $\delta = 2$, maximum coefficient magnitude $\alpha = 16.3384$, and base $b = 2$. Using Theorem 3 again, we find the $P \geq 18$ and $p \geq 8$ suffice.


 Figure 2. The network \mathcal{G}_2 .

However, these estimates for P and p are far from being tight. We first note that, along any path from source to terminal, there are only two non-terminal nodes that perform a linear combination. Thus, the “effective” depth is only 3 in this case. Furthermore, by using the exact values of the various $\alpha_{e \rightarrow e'}$, instead of the upper bound α , we can obtain a tighter bound on the maximal value passing over an edge, and the maximal error ϵ_{d-1} . In this case, after a simple processing by a computer, these give $P \geq 14$ and $p \geq 6$.

It follows that the network \mathcal{G}_2 , using the quasi-linear network coding system, is capable of transmitting source message of length 7 bits, using messages of length $14 + 6 = 20$ bits, i.e., with rate $7/20 > 1/3$. \square

Example 6. As a final example we bring the network \mathcal{G}_3 , shown in Figure 3, which is a combination of the networks \mathcal{G}_1 and \mathcal{G}_2 . This combination was shown to have no scalar linear solution in [3], though an ad-hoc non-linear solution was given.

We can apply the quasi-linear network-coding scheme to this network, and since the two sub-networks operate separately, get a solution with rate $7/20 > 1/3$. We can compare this with the best routing solution for this network which has a lower rate of $1/3$ (see [2]).

We are not forced in any way to use the quasi-linear scheme for the entire network. By combining a routing solution for \mathcal{G}_2 with rate $2/3$ (see [2]), with a quasi-linear solution for \mathcal{G}_1 with rate $\frac{n}{n+2}$, for any n , we can obtain an overall solution with rate $2/3$. \square

IV. CONCLUSION

In this paper we described quasi-linear network coding, which is a two-phase heuristic for designing vector non-linear network codes for non-multicast networks. In the first stage an approximate solution over the reals is found, and in the

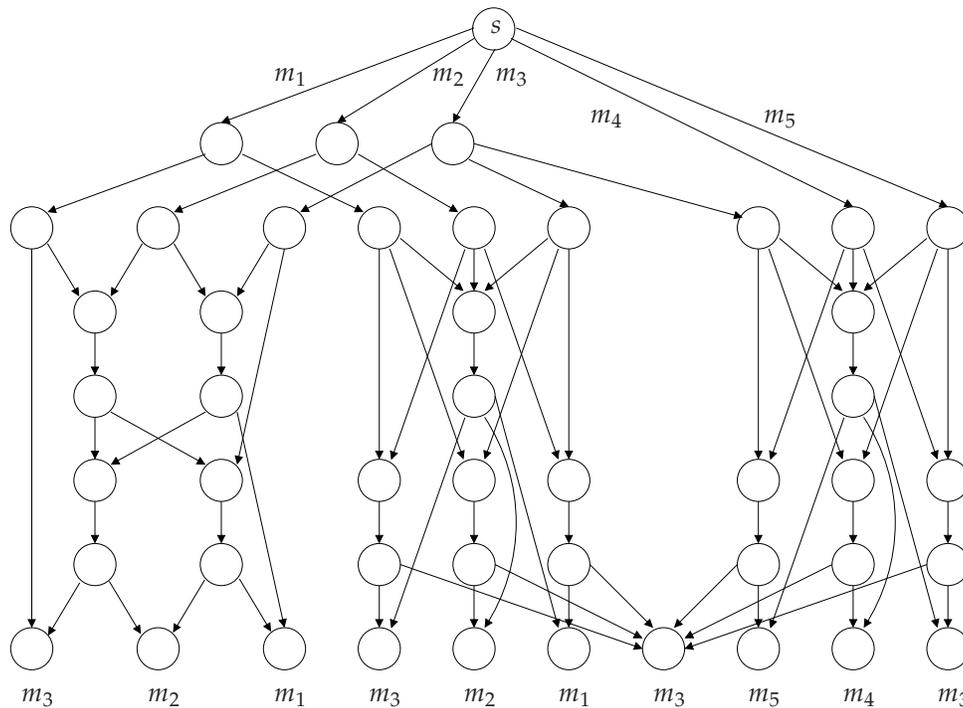


Figure 3. The network \mathcal{G}_3 .

second stage it is quantized to a fixed-point representation. We analyzed the sources for errors in the process and determined sufficient conditions for zero-error at the terminals. These conditions determine the rate of the solution.

We applied the method to the network presented in [3] to prove the insufficiency of linear network codes. While overall the rate was below that of the ad-hoc non-linear solution given in [3], our method is systematic, and it out-performs the routing capacity of the network.

Connections between the rate of the quasi-linear network code, and various parameters of the network, e.g., depth and incoming degree, were established. However, a crucial piece is still missing, and that is connecting the approximation factor γ with the network. This missing link will enable us to fully compare quasi-linear network codes with other coding techniques.

ACKNOWLEDGMENTS

The first author would like to thank the second author for hosting him at MIT during his sabbatical.

REFERENCES

- [1] R. Ahlswede, N. Cai, S. Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.
- [2] J. Cannons, R. Dougherty, C. Freiling, and K. Zeger, "Network routing capacity," *IEEE Trans. Inform. Theory*, vol. 52, no. 3, pp. 777–788, Mar. 2006.
- [3] R. Dougherty, C. Freiling, and K. Zeger, "Insufficiency of linear coding in network information flow," *IEEE Trans. Inform. Theory*, vol. 51, no. 8, pp. 2745–2759, Aug. 2005.
- [4] —, "Networks, matroids, and non-shannon information inequalities," *IEEE Trans. Inform. Theory*, vol. 53, no. 6, pp. 1949–1969, Jun. 2007.
- [5] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Trans. Inform. Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.
- [6] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. M. G. M. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Inform. Theory*, vol. 51, no. 6, pp. 1973–1982, Jun. 2005.
- [7] A. Kim and M. Médard, "Scalar-linear solvability of matroidal networks associated with representable matroids," *Proc. Int. Symp. Turbo Codes and Iterative Information Processing*, 2010.
- [8] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. on Networking*, vol. 11, no. 5, pp. 782–795, Oct. 2003.
- [9] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inform. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.
- [10] M. Médard, M. Effros, D. Karger, and T. Ho, "On coding for non-multicast networks," in *Proceedings of the Annual Allerton Conference on Communication Control and Computing*, vol. 41, no. 1, 2003, pp. 21–29.
- [11] V. T. Muralidharan and B. S. Rajan, "Linear network coding, linear index coding and representable discrete polymatroids," *CoRR*, vol. abs/1306.1157, 2013.
- [12] A. Rasala Lehman and E. Lehman, "Complexity classification of network information flow problems," in *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2004, pp. 142–150.
- [13] S. Riis, "Linear versus non-linear boolean functions in network flow," in *38th Annual Conference on Information Science and Systems (CISS)*, Princeton, NJ, 2004.
- [14] S. Shintre, S. Katti, S. Jaggi, B. K. Dey, D. Katabi, and M. Médard, "Real and complex network codes: promises and challenges," in *Proceedings of the 2008 Workshop on Network Coding, Theory, and Applications (NetCod2008)*, Hong Kong, Jan. 2008.