The Local Mixing Problem

Yunnan Wu, Jitendra Padhye, Ranveer Chandra, Venkat Padmanabhan, Philip A. Chou
Microsoft Research, One Microsoft Way, Redmond, WA, 98052.

***Abstract*—The presence of broadcast channels renders network coding particularly useful. For example, a single broadcast transmission of a proper mixture packet may simultaneously present useful information to multiple nodes. Motivated by the practical application of network coding in wireless networks, this paper formulates a *local mixing* problem: A source node has a set of mutually independent sources; each source is available at a subset of neighbors and needs to be transferred to another subset of neighbors; the problem is to characterize the *admissible rate region*, i.e., the set of channel rate allocations that can fulfill the traffic demand. This paper establishes that under the constraint that each neighbor decodes only from received symbols that are functions of sources that the node can eventually recover, the admissible rate region can be characterized by a set of linear constraints and linear mixing is optimal.**

I. Introduction

Network coding refers to a scheme where a node is allowed to generate output data by mixing (i.e., computing certain functions of) its received data. This extends conventional routing, which allows a node to only forward its received data. This concept was first introduced by Ahlswede et al. [1] in the context of multicasting data in a network of lossless links.

Network coding is particularly useful in a broadcast medium. For example, consider three wireless nodes v_0, v_1, v_2 illustrated in Figure 1, where nodes v_1 and v_2 are both within the communication range of node v_0. Suppose node v_1 has packet $\boldsymbol{x}_1$, node v_2 has packet $\boldsymbol{x}_2$, and node v_0 has packets $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$. Suppose further that node v_1 needs packet $\boldsymbol{x}_2$ and node v_2 needs packet $\boldsymbol{x}_1$. A single transmission of a packet $\boldsymbol{x}_1 + \boldsymbol{x}_2$ (where '+' stands for the bit-wise XOR of the two packets) by node v_0 achieves two purposes: It allows node v_1 to recover $\boldsymbol{x}_2$ and node v_2 to recover $\boldsymbol{x}_1$. This technique was termed *physical piggybacking* by Wu et al. [2] because the two packets are combined into one, without even increasing the size of the packet.

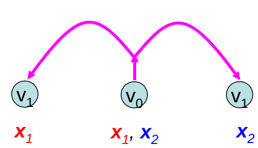

Fig. 1. Physical piggybacking: A single broadcast transmission of $\boldsymbol{x}_1 + \boldsymbol{x}_2$ presents $\boldsymbol{x}_2$ to node v_1 who knows $\boldsymbol{x}_1$ and $\boldsymbol{x}_1$ to node v_2 who knows $\boldsymbol{x}_2$.

Wu et al. [2] demonstrated that packet exchanges can naturally benefit from physical piggybacking. More recently, Katti et al. [3] presented a framework for taking advantage of physical piggybacking to improve the efficiency of unicasting in multi-hop wireless networks. In their framework, each node snoops on the medium and buffers packets it heard. A node

Neighbor \ Field	Has	Wants
v_1	x_1, x_3	x_4
v_2	x_1, x_4	x_3
v_3	x_2, x_3, x_4	x_1, x_5
v_4	x_1, x_5	x_2

Fig. 2. A table representing the traffic demand. For each neighbor v, $v.Has$ denotes the set of sources that v has; $v.Wants$ denotes the set of sources that v wants to receive.

also informs its neighbors which packets they overheard. This allows nodes to know (roughly) what packets are available at each neighbor and then perform mixing to efficiently use the channel.

In such a system, a critical issue is how to optimize the formation of mixture packets so as to most efficiently use the channel resource. Figure 2 gives one example problem scenario. The table lists the traffic demand. For this example, one possible solution is to broadcast $\boldsymbol{x}_1 + \boldsymbol{x}_3 + \boldsymbol{x}_4$ to $\{v_1, v_2, v_3\}$ and $\boldsymbol{x}_2 + \boldsymbol{x}_5$ to $\{v_3, v_4\}$.

In this paper we formulate a mathematical abstraction of such an issue, which is called the *local mixing* problem. There is a source node v_0, who has a set of neighbors V. The source node v_0 has M mutually independent sources, $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M$. The rate of source $\boldsymbol{x}_m$ is r_m. Each source $\boldsymbol{x}_m$ is initially available at a set of neighbors $H_m \subseteq V$ and is needed by a set of neighbors $W_m \subseteq V$. Hence the traffic demand can be characterized by M tuples $\{\langle H_m, W_m, r_m \rangle\}_{m=1}^{M}$.

The traffic demand needs to be fulfilled by the channel resources. We abstract out the channel coding issue and focus on the mixing (network coding) issue. This amounts to a separate treatment of channel coding and network coding, where the channel coding is concerned about converting the noisy channel into near-lossless bits and the network coding is concerned about using the lossless bits to fulfill the traffic demand. Although this separation approach could be suboptimal than a joint channel and network coding approach, the separation approach is closer to engineering practice and it allows us to focus on mixing without referring to channel properties. Specifically, the physical layer is abstracted by a capacity region $\mathcal{C}$ consisting of channel rate vectors $\boldsymbol{c}$ of length $2^{|V|} - 1$. In the case there are three neighbors, v_1, v_2, v_3, each channel rate vector is of the form

$$\boldsymbol{c} \triangleq [c_{\{1\}}, c_{\{2\}}, c_{\{3\}}, c_{\{1,2\}}, c_{\{1,3\}}, c_{\{2,3\}}, c_{\{1,2,3\}}], \quad (1)$$

meaning that the channel can simultaneously provide rate $c_{\{1\}}$ to v_1, rate $c_{\{1,2\}}$ to $\{v_1, v_2\}$, etc. In other words, we are modelling the physical layer as $2^{|V|} - 1$ lossless channels indexed by the set of receivers; the channel for a subset $Q \subseteq V$ can transfer information reliably from v_0 to Q at rate c_Q.

Therefore, the problem is to characterize the *admissible rate region*, i.e., the set of channel rate allocations $\boldsymbol{c}$ that can fulfill the traffic demand $\{\langle H_m, W_m, r_m \rangle\}_{m=1}^{M}$.

The main result of this paper is that under the constraint that each neighbor decodes only from received symbols that are functions of sources that the node can eventually recover, the admissible rate region can be characterized by a set of linear constraints and linear mixing is optimal.

II. Optimal Mixing Assuming Pollution-Free Decoding

Consider traffic demand $\{\langle H_m, W_m, r_m \rangle\}_{m=1}^{M}$. For the rest of the paper, we use a different representation of the traffic demand. We can classify the sources into several types: A source characterized by $\langle H_m, W_m, r_m \rangle$ is said to be of type $\langle H_m, W_m \rangle$. Let Φ denote the set of all possible source types. For example, if there are two neighbors $V = \{v_1, v_2\}$, then

$$\Phi = \{\langle \emptyset, \{1\}\rangle, \langle \emptyset, \{2\}\rangle, \langle \emptyset, \{1,2\}\rangle, \langle \{1\}, \{2\}\rangle, \langle \{2\}, \{1\}\rangle\} .$$

We represent the original traffic demand $\{\langle H_m, W_m, r_m \rangle\}_{m=1}^{M}$ as a length-$|\Phi|$ vector $\boldsymbol{r}$, whose entry r_ϕ is the sum rate of sources of type $\phi \in \Phi$. We then treat as if there are $|\Phi|$ sources, with rate specified by $\boldsymbol{r}$.

Any solution for the problem sends out a sequence of output symbols on the broadcast channels. Each output symbol is a function of some of the sources $\{\boldsymbol{x}_\phi\}$.

Each neighbor $v \in V$ receives a subset of the output symbols and needs to recover its needed sources using the received output symbols, and the source symbols that it initially has. For each neighbor v, denote the set of sources that v has or wants by

$$\underline{\boldsymbol{x}}_v \triangleq \{\boldsymbol{x}_\phi \in \Phi : \ v \in \phi_H \cup \phi_W\}. \tag{2}$$

Here we write $\phi = \langle \phi_H, \phi_W \rangle$ as a convention. Sources outside $\underline{\boldsymbol{x}}_v$ are said to be "pollution sources" for v. If a pollution source is involved in generating an output symbol Y, then we say Y is a "polluted" symbol from v's point of view. In this section we impose the following constraint and look for optimal solutions under this constraint.

Constraint 1 ("Decoding from Unpolluted Data Only"): Each node $v \in V$ decodes its needed information using only the sources it initially has, and its received output symbols that are functions of sources in $\underline{\boldsymbol{x}}_v$. In other words, each node is required to ignore any polluted symbols that it receives.

Theorem 1 (Two-Stage Assignment): Assuming pollution-free decoding (i.e., under Constraint 1), the traffic demand $\boldsymbol{r}$ can be fulfilled by channel rate vector $\boldsymbol{c}$

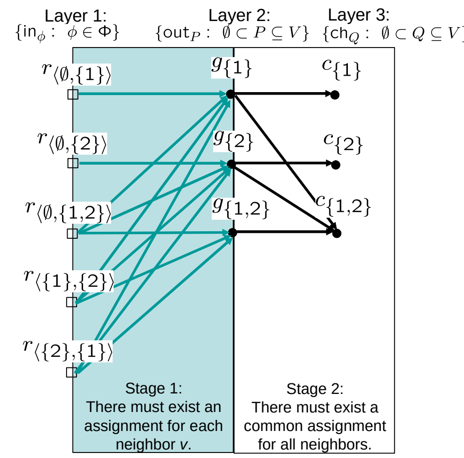

Fig. 3. Illustration of the linear constraints in Theorem 1.

if and only if there exists $\{f_{\phi P}^{(v)}, g_P, \sigma_{PQ}\}$ that satisfy

$$\sum_{P} f_{\phi P}^{(v)} = r_\phi, \ \ \forall v \in V, \ \forall \phi : v \in \phi_W \tag{3}$$

$$\sum_{\phi} f_{\phi P}^{(v)} \le g_P, \ \ \forall v \in V, \ \forall P : \{v\} \subseteq P \subseteq V, \tag{4}$$

$$f_{\phi P}^{(v)} \ge 0, \qquad \forall P, \phi, v : \{v\} \subseteq P \subseteq \phi_H \cup \phi_W \subseteq V, \tag{5}$$

$$\sum_{Q} \sigma_{PQ} = g_P, \ \forall P : \emptyset \subset P \subseteq V \tag{6}$$

$$\sum_{P} \sigma_{PQ} \le c_Q, \ \forall Q : \emptyset \subset Q \subseteq V \tag{7}$$

$$\sigma_{PQ} \ge 0, \qquad \forall P, Q : \ \emptyset \subset P \subseteq Q \subseteq V. \tag{8}$$

Here variable $f_{\phi P}^{(v)}$ exists only for $\{v\} \subseteq P \subseteq \phi_H \cup \phi_W \subseteq V$; variable σ_{PQ} exists only for $\emptyset \subset P \subseteq Q \subseteq V$.

Furthermore, if the above linear system of constraints has a feasible solution, then the traffic demand $\boldsymbol{r}$ can be fulfilled by $\boldsymbol{c}$ via linear coding.

We now explain the linear constraints in Theorem 1. We introduce a directed graph $\boldsymbol{Z}$, as illustrated by Fig. 3. There are three layers of vertices. The nodes in layer 1, $\{in_\phi : \phi \in \Phi\}$, model the $|\Phi|$ sources; the node in_ϕ has an associated (source) traffic r_ϕ. The nodes in layer 2, $\{out_P : \emptyset \subset P \subseteq V\}$, model the $2^{|V|} - 1$ "output buffers" that hold output symbols (their meaning will be explained shortly); the buffer out_P has an associated rate g_P, representing the total rate of its output symbols. The nodes in layer 3, $\{\mathrm{ch}_Q : \emptyset \subset Q \subseteq V\}$, model the $2^{|V|} - 1$ physical channels; the channel ch_Q has an associated capacity c_Q.

A source node in_ϕ has an outgoing edge to an output buffer out_P if $\phi_H \cup \phi_W \supseteq P$; such an edge is denoted by ϕP. An output buffer out_P has an outgoing edge to channel ch_Q if $P \subseteq Q$; such an edge is denoted by PQ. We use the name "stage 1" (resp. "stage 2") to refer to the subgraph of $\boldsymbol{Z}$ induced by layer-1 and layer-2 nodes (resp. layer-2 and layer-3 nodes).

In Theorem 1, Constraints (3)-(5) correspond to an traffic assignment in stage 1 *for each neighbor* v, where the sources that v wants are assigned to the output buffers $\{out_P\}$ that v receives. Constraints (6)-(8) correspond to an traffic assignment in stage 2, where the traffic held by buffers $\{out_P\}$ are assigned to the channels $\{\mathrm{ch}_Q\}$.

Theorem 1 can now be alternatively stated as follows. Assuming pollution-free decoding, traffic demand $\boldsymbol{r}$ can be fulfilled by channel resource $\boldsymbol{c}$ if and only if there exists a valid traffic assignment $\{f^{(v)}_{\phi P}\}$ in stage 1 *for each individual neighbor* v and a common valid traffic assignment $\{\sigma_{PQ}\}$ in stage 2 *for all neighbors*.

A. Proof of Necessity

Any solution $\mathcal{S}$ corresponds to an assignment of output symbols to the $2^{|V|}-1$ channels, such that the total rate in each channel Q does not exceed the channel capacity c_Q.

Consider assigning each output symbol into one of $(2^{|V|}-1)$ output buffers $\{out_P : \emptyset \subset P \subseteq V\}$ as follows. For each output symbol Y, we put it into buffer out_P, where P is the set of neighbors that uses Y for decoding. After such a classification process, let g_P denote the sum rate of output symbols in buffer out_P; collectively, let $\boldsymbol{g}$ be the length-$(2^{|V|}-1)$ vector obtained by concatenating $\{g_P\}$ together. Note that if a symbol is transmitted by channel ch_Q and used by P for decoding, then $P \subseteq Q$. Let σ_{PQ} be the sum rate of symbols that are transmitted by channel ch_Q and used by P for decoding. Then (6)-(8) must be satisfied.

Now consider an arbitrary neighbor v. It has access to the symbols in buffers $\{out_P : P \ni v\}$, and the sources $\{\boldsymbol{x}_{\phi'} : \phi_H \ni v\}$. It needs to recover sources $\{\boldsymbol{x}_\phi : \phi_W \ni v\}$. Due to Constraint 1, each symbol in out_P can only be a function of $\{\boldsymbol{x}_\phi : \phi_H \cup \phi_W \supseteq P\}$.

The coding for node v can be viewed as a single source network coding problem in a network, illustrated by Fig. 4. In $\boldsymbol{Z}$, add a virtual source node s_v that has an outgoing edge with capacity r_ϕ to each source $\boldsymbol{x}_\phi$ that v wants. In addition, add a virtual destination node t_v that has an incoming edge with infinite capacity from each source $\boldsymbol{x}_{\phi'}$ that v has, and an incoming edge with capacity g_P from each output buffer out_P with $P \ni v$. Denote the resulting graph by $\boldsymbol{Z}_v$. Then any solution $\mathcal{S}$ for the original local mixing problem maps into a solution in $\boldsymbol{Z}_v$ for the problem of transferring the sources $\underline{\boldsymbol{x}}_v$ from s_v to t_v. From [1], there must exist an s_v–t_v flow in $\boldsymbol{Z}_v$ that provides rate $\sum_{\phi:\phi_H\cup\phi_W\ni v} r_\phi$. Such flow condition is equivalent to the constraints (3)-(5).

Therefore the constraints (3)-(8) hold for any solution $\mathcal{S}$.

B. Proof of Sufficiency

We first explain the basic idea underlying the proof. We associate a tag with each output symbol; the value of the tag is a subset of V, which indicates the set of neighbors "interested" in this symbol. If a symbol Y is tagged with $U \subseteq V$, then Y must not be pollution to any neighbor in U; i.e.,

$$Y = f(\boldsymbol{x}_\phi : \phi_H \cup \phi_W \supseteq U), \text{ for some function } f. \tag{9}$$

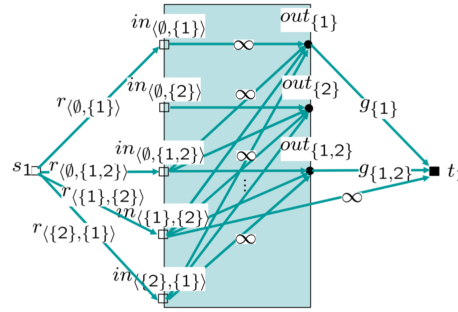

Fig. 4. The graph $\boldsymbol{Z}_v$.

We maintain $2^{|V|}-1$ tagged buffers $\{out_P : \emptyset \subset P \subseteq V\}$; the buffer out_P holds output symbols tagged with P. The output symbols in out_P are generated by mixing the compatible sources $\{\boldsymbol{x}_\phi : \phi_H \cup \phi_W \supseteq P\}$; this is modelled by the connection structure in stage 1 of $\boldsymbol{Z}$.

Proof of Sufficiency: Consider an arbitrary feasible solution, $\{f^{(v)}_{\phi P}, g_P, \sigma_{PQ}\}$, that satisfies the linear system of constraints (3)-(8). Without essential loss of generality, we assume all the variables and constants in (3)-(8) are integer.

In $\boldsymbol{Z}$, replace each edge ϕP in stage 1 by r_ϕ parallel unit-capacity edges from in_ϕ to out_P, and each edge PQ in stage 2 by σ_{PQ} parallel unit-capacity edges, $\{PQ_i : i = 1, \ldots, \sigma_{PQ}\}$, from out_P to ch_Q. Denote the resulting graph by Z. We next show that there exists a linear network coding assignment in Z that enables all neighbors to recover its needed information. Such a linear network coding assignment corresponds to a feasible solution that fulfills the traffic demand using channel resource $\boldsymbol{c}$. The proof is based on the algebraic framework introduced by Koetter and Medard [4].

Let $\mathbb{F}$ denote the operating finite field. Each unit-capacity edge in Z can carry one symbol from $\mathbb{F}$. The r_ϕ parallel edges from in_ϕ to out_P carry the r_ϕ source symbols, $x_{\phi 1}, \ldots, x_{\phi r_\phi}$, respectively. Let $\boldsymbol{y}_P$ denote the vector (of length g_P) formed by the set of all output symbols generated by buffer out_P. Then in a linear coding assignment, $\boldsymbol{y}_P$ is a linear transformation of the sources observed by buffer out_P:

$$\boldsymbol{y}_P = \sum_{\phi:\ \phi P \in E(\boldsymbol{Z})} \boldsymbol{w}_{\phi P} \boldsymbol{x}_\phi \tag{10}$$

The elements in $\boldsymbol{w}_{\phi P}$ are said to be "mixing coefficients". Let $\boldsymbol{w}$ be a vector that includes every $\boldsymbol{w}_{\phi P}$ for $\phi P \in E(\boldsymbol{Z})$. Hence a linear network coding assignment is specified by an assignment of $\boldsymbol{w}$, from $\mathbb{F}^{|\boldsymbol{w}|}$ (here $|\boldsymbol{a}|$ denotes the length of vector $\boldsymbol{a}$).

Consider an arbitrary neighbor v. We treat $\underline{\boldsymbol{x}}_v$ defined in (2) as a vector formed by concatenating the sources that v has or wants. It has access to all the channels $\{ch_Q : Q \ni v\}$. Since all symbols generated by out_P are assigned to channels $\{ch_Q : P \subseteq Q\}$ in stage 2, v has access to the output symbols $\{\boldsymbol{y}_P : P \ni v\}$. In addition, v has access to the source symbols $\{\boldsymbol{x}_\phi : \phi_H \ni v\}$. Let $\underline{\boldsymbol{y}}_v$ be a vector formed by $\{\boldsymbol{y}_P : P \ni v\}$ and $\{\boldsymbol{x}_\phi : \phi_H \ni v\}$. For each $P \ni v$, note that $\boldsymbol{y}_P$ is a function of only $\underline{\boldsymbol{x}}_v$, due to the connection structure of $\boldsymbol{Z}$. Thus $\underline{\boldsymbol{y}}_v$ is a linear function of $\underline{\boldsymbol{x}}_v$:

$$\underline{\boldsymbol{y}}_v = Q_v \underline{\boldsymbol{x}}_v, \tag{11}$$

where each element of Q_v is either a binary constant or a mixture coefficient in $\boldsymbol{w}$.

The needed source information $\underline{\boldsymbol{x}}_v$ can be decoded from $\underline{\boldsymbol{y}}_v$ if and only if Q_v has full column rank, which holds if and only if there exists a matrix P_v of size $|\underline{\boldsymbol{x}}_v| \times |\underline{\boldsymbol{y}}_v|$ such that $\det(P_v Q_v) \neq 0$. Note that $\{\boldsymbol{f}^{(v)}_{\phi P}\}$ is a valid assignment of the sources traffic $\{r_\phi : \phi_W \ni v\}$ to the buffers $\{out_P : P \ni v\}$. Thus, there exists an assignment of $\boldsymbol{w}$ and P_v, with each element being either 0 or 1, such that $\det(P_v Q_v) \neq 0$.

Now we consider the recoverability of all destination nodes. There exists a linear coding solution such that each destination can recover its needed information if and only if there exists an assignment of $\boldsymbol{w}$ and $\{P_v : v \in V\}$ that satisfies

$$\prod_{v \in V} \det(P_v Q_v) \neq 0. \tag{12}$$

Following [4], the quantity $\prod_{v \in V} \det(P_v Q_v)$ can be viewed in two ways: (i) as a polynomial in terms of the variables $\boldsymbol{w}$ and $\{P_v : v \in V\}$ with coefficients in $\mathbb{F}$, (ii) as a number in $\mathbb{F}$ for given $\boldsymbol{w}$ and $\{P_v : v \in V\}$. Earlier we have established that for each destination v alone, there exists an assignment of the elements in $\boldsymbol{w}$ and P_v from $\mathbb{F} = GF(2)$ such that $\det(P_v Q_v) \neq 0$. This implies that for each destination v, $\det(P_v Q_v)$ is a nonzero polynomial. Therefore, $\prod_{v \in V} \det(P_v Q_v)$ is a nonzero polynomial in terms of the variables $\boldsymbol{w}$ and $\{P_v : v \in V\}$ with coefficients in $\mathbb{F}$. It is known that for a non-zero polynomial defined over a sufficiently large finite field $\mathbb{F}$, it must evaluate to a nonzero value at a certain point. This establishes the existence of $\boldsymbol{w}$ and $\{P_v : v \in V\}$ such that all needed symbols are recovered simultaneously at all destinations.

In the above we have shown the existence of a linear coding solution for a sufficiently large field $\mathbb{F}$. In fact, if we choose each mixing coefficient in $\boldsymbol{w}$ and each element of $\{P_v : v \in V\}$ uniformly and independently from $\mathbb{F}$, then the probability that all needed symbols are recovered approaches 1 as the field size $|\mathbb{F}|$ approaches infinity; this can be established via the Schwartz-Zippel Theorem (e.g., [5], according to [6]).

III. Optimal Mixing Under a Relaxed Constraint

In this section we show it is possible to characterize the optimal mixing under Constraint 2, which is a relaxed version of Constraint 1.

Constraint 2 ("Extended Pollution-Free Decoding): If node v uses an output symbol Y that is a function of a subset of source symbols $\mathcal{X}$, for decoding, then node v must recover all source symbols in $\mathcal{X}$.

Constraint 2 relaxes Constraint 1 by allowing a node v to treat some unwanted source symbols as wanted symbols. Equivalently, some source symbols now have more recipients. This results in more symbols to be recovered; the benefit is that some polluted symbols under Constraint 1 now become symbols that can be used by v for decoding.

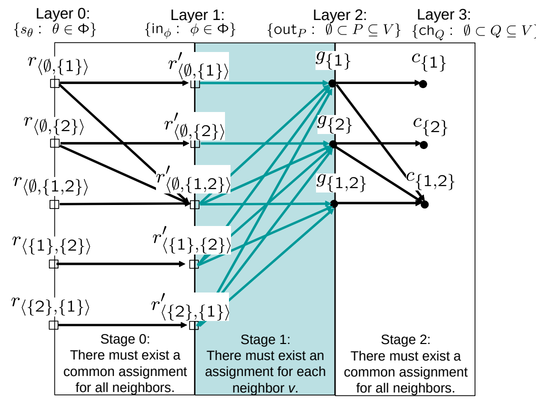

Fig. 5. Illustration of the linear constraints in Theorem 2.

The degree of freedom in adding recipients can be modelled by an additional assignment stage, as illustrated by Fig. 5 for the case $V = \{v_1, v_2\}$. We can treat a part of the original source $\boldsymbol{x}_\theta$ as source symbols of a source $\boldsymbol{x}'_\phi$ with $\theta_H = \phi_H$ and $\theta_W \subseteq \phi_W$. Let $\tau_{\theta\phi}$ denote the amount of traffic which is originally of type θ and converted to type ϕ. Then (13)-(15) hold. After such an assignment stage, the original traffic demand $\boldsymbol{r}$ is now converted into a traffic demand $\boldsymbol{r}'$. Therefore, we have the following extension of Theorem 2.

Theorem 2 (Three-Stage Assignment): Under Constraint 2, the traffic demand $\boldsymbol{r}$ can be fulfilled by channel rate vector $\boldsymbol{c}$ if and only if there exists $\{\tau_{\theta\phi}, r'_\phi, f^{(v)}_{\phi P}, g_P, \sigma_{PQ}\}$ that satisfy

$$\sum_\phi \tau_{\theta\phi} = r_\theta, \ \forall\theta \in \Phi, \tag{13}$$

$$\sum_\theta \tau_{\theta\phi} = r'_\phi, \ \forall\phi \in \Phi, \tag{14}$$

$$\tau_{\theta\phi} \geq 0, \ \ \forall\theta, \phi \in \Phi : \theta_H = \phi_H, \theta_W \subseteq \phi_W. \tag{15}$$

$$\sum_P f^{(v)}_{\phi P} = r'_\phi, \ \ \forall v \in V, \ \forall\phi : v \in \phi_W \tag{16}$$

$$\sum_\phi f^{(v)}_{\phi P} \leq g_P, \ \ \forall v \in V, \ \forall P : \{v\} \subseteq P \subseteq V, \tag{17}$$

$$f^{(v)}_{\phi P} \geq 0, \qquad \forall P, \phi, v : \{v\} \subseteq P \subseteq \phi_H \cup \phi_W \subseteq V, \tag{18}$$

$$\sum_Q \sigma_{PQ} = g_P, \ \forall P : \emptyset \subset P \subseteq V \tag{19}$$

$$\sum_P \sigma_{PQ} \leq c_Q, \ \forall Q : \emptyset \subset Q \subseteq V \tag{20}$$

$$\sigma_{PQ} \geq 0, \qquad \forall P, Q : \ \emptyset \subset P \subseteq Q \subseteq V. \tag{21}$$

Here variable $\tau_{\theta\phi}$ exists only for $\theta, \phi \in \Phi : \theta_H = \phi_H, \theta_W \subseteq \phi_W$; variable $f^{(v)}_{\phi P}$ exists only for $\{v\} \subseteq P \subseteq \phi_H \cup \phi_W \subseteq V$; variable σ_{PQ} exists only for $\emptyset \subset P \subseteq Q \subseteq V$.

Furthermore, if the above linear system of constraints has a feasible solution, then the traffic demand $\boldsymbol{r}$ can be fulfilled by $\boldsymbol{c}$ via linear coding.

IV. Relation with Multi-Source Multicasting

The local mixing problem can be viewed as a special case of the general multi-source multicasting problem (see, e.g., [7]). In the multi-source multicasting problem, there are a set of independent sources; each source is to be multicast through a network of lossless channels to a set of destination nodes.

We now show how to interpret the local mixing problem as a multi-source multicasting problem. We introduce a graph consisting of four layers of nodes, as illustrated by Figure 6. The first layer consists of one node for each source type $\phi \in \Phi$. The second layer consists of a single node v_0. The third layer consists of one node for each nonempty subset of V. The fourth layer consists of one node for each neighbor in V. There is an edge with infinite capacity from every source X_ϕ to v_0, representing that v_0 has access to all these sources. Each source X_ϕ has edges with infinite capacity pointing to neighbors that have the source. The use of the channel for broadcasting to a set $Q \subseteq V$ is represented by a tree-like structure: There is an edge with capacity c_Q from v_0 to the node Q in the third layer and there is an edge with infinite capacity from Q to each neighbor in Q.

Therefore the local mixing problem is a special case of the multi-source multicasting problem. As a result, known theoretical results about the multi-source multicasting problem can be applied. In particular, the bounding techniques by Song et al. [8] (see also Chapter 15 of [7]) can be applied. We now explain the results.

Introduce random variables $\{X_\phi : \phi \in \Phi\}$ and $\{Y_Q : \emptyset \subset Q \subseteq V\}$. Let $\mathcal{N}$ denote the union of these two sets. Let $n \triangleq |\mathcal{N}| = |\Phi| + 2^{|V|} - 1$. Let $\mathcal{H}_n$ denote the $(2^n - 1)$-dimensional Euclidean space with the coordinates labelled by $\{h_A : \emptyset \subset A \subseteq \mathcal{N}\}$. A column vector $\boldsymbol{h} \in \mathcal{H}_n$ is called *entropic* if for a certain joint distribution of random variables in $\mathcal{N}$, h_A is equal to the joint entropy of the random variables A, for every coordinate $A \subseteq \mathcal{N}$. Then, the set of all entropic vectors is called the *entropy space* Γ_n^* of n random variables:

$$\Gamma_n^* \triangleq \{\boldsymbol{h} \in \mathcal{H}_n : \boldsymbol{h} \text{ is entropic}\}. \tag{22}$$

Let $\overline{\Gamma_n^*}$ denote the closure of Γ_n^*. For ease in notation, define $h_{A|B} \triangleq h_{A\cup B} - h_B$.

Specifying the inner bound of [8] and after some manipulations, we see that $\boldsymbol{c}$ can fulfill traffic demand $\boldsymbol{r}$ if there exist $\boldsymbol{h} \in \Gamma_n^*$ that satisfies the following conditions:

$$h_{(X_\phi:\phi\in\Phi)} = \sum_{\phi\in\Phi} h_{X_\phi}, \tag{23}$$

$$h_{Y_Q|(X_\phi:\,\phi\in\Phi)} = 0, \tag{24}$$

$$h_{(X_\phi:v\in\phi_W)|(X_\phi:v\in\phi_H),(Y_Q:\,v\in Q)} = 0,\ \forall v \in V, \tag{25}$$

$$c_Q > h_{Y_Q}, \quad \forall Q : \emptyset \subset Q \subseteq V, \tag{26}$$

$$h_{X_\phi} > r_\phi, \quad \forall \phi \in \Phi. \tag{27}$$

These inequalities can be interpreted as follows. Here (23) says that the sources are mutually independent; (24) says that the output symbols has to be a function of the inputs; (25) says that each node must be able to recover all its wanted sources;

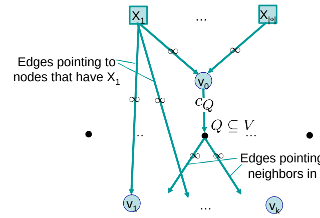

Fig. 6. Interpretation of the local mixing problem as a multi-source problem.

(26) says that the entropy of the output Y_Q must be strictly less than the total available channel bit-rate; (27) says that the entropy of X_ϕ must be strictly less than the information rate of source ϕ.

The outer bound of [8] is applicable. Furthermore, following the derivations of [8] (or Chapter 15 of [7]), an outer bound customized for the current context can be obtained. The result is that if $\boldsymbol{c}$ can fulfill traffic demand $\boldsymbol{r}$, then there exists $\boldsymbol{h} \in \overline{\Gamma_n^*}$ that satisfies the linear constraints obtained by replacing '$>$' in (26)(27) by '$\geq$'. Moreover, a relaxed bound can be obtained by replacing $\overline{\Gamma_n^*}$ by its outer bound Γ_n, which is the set of nonnegative vectors $\boldsymbol{h} \in \mathcal{H}_n$ that satisfies all Shannon-type inequalities; see Chapter 15 of [7] for details.

V. Conclusion

We introduced the local mixing problem and presented constructive results. Theorem 1 states that if a system of linear constraints has a feasible solution, then the traffic demand can be fulfilled by the given channel resource via (random) linear coding. This result is established via two key constructive techniques: (i) organizing data into different categories based on the interested neighbors, (ii) systematically exploring the degree of freedom in adding recipients. The local mixing problem is a special case of the multi-source multicasting problem: It does not involve multi-hop relaying. The techniques in this paper have been extended to the general multi-source multicasting problem in [9].

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, July 2000.

[2] Y. Wu, P. A. Chou, and S.-Y. Kung, "Information exchange in wireless networks with network coding and physical-layer broadcast," in *Proc. 39th Annual Conf. Inform. Sci. and Systems (CISS)*, Mar. 2005.

[3] S. Katti, D. Katabi, W. Hu, H. Rahul, and M. Medard, "Practical network coding for wireless environments," in *Proc. 43rd Annual Allerton Conf. Comm., Ctrl., and Comp.* IEEE, Sept. 2005.

[4] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 782–795, Oct. 2003.

[5] R. Motwani and P. Raghavan, *Randomized Algorithms*. Cambridge University Press, 1995.

[6] 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*, 2004, submitted. http://web.mit.edu/trace/www/.

[7] R. W. Yeung, *A First Course in Information Theory*. Kluwer Academic/Plenum Publishers, 2002.

[8] L. Song, R. W. Yeung, and N. Cai, "Zero-error network coding for acyclic networks," *IEEE Trans. Inform. Theory*, vol. 49, no. 12, pp. 3129–3139, Dec. 2003.

[9] Y. Wu, "On constructive multi-source network coding," in *Submitted to Int'l Symp. Inform. Theory (ISIT 2006)*. IEEE, Jan. 2006.