

Applications of Sheaf Cohomology and Exact Sequences to Network Coding

Robert Ghrist*

Yasuaki Hiraoka†

Abstract

This paper introduces a novel application of sheaf cohomology — a mathematical tool for collating local algebraic data into global structures — to network coding problems. After recalling the definition of sheaves, we introduce NETWORK CODING SHEAVES for a general multi-source network coding scenario, and consider various forms of sheaf cohomologies. The main theorem states that 0-th network coding sheaf cohomology is equivalent to information flows for the network coding. This theorem is then applied to several practical problems in network codings such as maxflow bounds, global extendability, network robustness, and data merging; these applications all follow from exact sequences of sheaf cohomology.

1 Introduction

This paper introduces new tools for the analysis of data flows over networks. We focus on (linear) network coding [2, 15], an important class of problems with numerous applications to error correction, optimal throughput, network security, and distribution [3]. Network coding is one of a host of problems in data analysis and management that require an understanding of local-to-global transitions. The novel tools we present in this paper are based on sheaf theory [1, 4, 8, 10].

Sheaf theory was invented in the mid 1940s as a branch of algebraic topology to organize local structures on topological spaces [10, Introduction]. Via its successes in several complex variables and algebraic geometry, sheaves are now indispensable in modern mathematics. However, despite its prowess in dealing with local-to-global transitions, sheaf theory has rarely if ever found concrete applications to problems in science or engineering. The few exceptions in the literature (e.g., [5, 14]) have dealt with applications to logic/semantics in Computer Science and use categorical properties of sheaves to organize local data.

This paper introduces the following principal ideas:

1. SHEAVES are an excellent tool for organizing network information flows;
2. SHEAF COHOMOLOGY yields global characterizations of networks with coding;
3. EXACT SEQUENCES allow for easy manipulation and computation of the above.

*University of Pennsylvania: gchrist@math.upenn.edu

†Hiroshima University/JST: hiraok@hiroshima-u.ac.jp

This note is part of a larger program to introduce sheaf cohomology as a means to organize data over networks: the recent complementary work of Robinson is the effect of these investigations to organize data over quantum graphs [12] and logic circuits [13].

The paper is organized as follows. In Section 2, we overview the fundamentals of sheaf theory, including sheaf cohomology and basic operations. This section introduces the NETWORK CODING SHEAF (NC sheaf for short). We then demonstrate the information-theoretical meaning of NC sheaf cohomology. In Section 3, NC sheaf cohomology is applied to certain practical problems: maxflow bounds, global extendability, network robustness, and data merging. The techniques of this section are standard long exact sequences in homological algebra.

2 Sheaf Formulation of Network Codings

2.1 Definition of Sheaves

The reader may consult [1, 4, 8, 10] for general discussions on sheaf theory. Let X be a topological space (e.g., network, as a 1-d cell complex) and \mathcal{R} be a commutative ring.

Definition 1 (Presheaf). A PRESHEAF F on X consists of the following data:

1. an \mathcal{R} -module $F(U)$ for each open subset $U \subset X$;
2. an \mathcal{R} -linear map $\rho_{VU} : F(U) \rightarrow F(V)$ for each pair $V \subset U \subset X$.

These data satisfy the following conditions:

$$\rho_{UU} = \text{ID}_U, \rho_{WV} \circ \rho_{VU} = \rho_{WU} \text{ for } W \subset V \subset U,$$

where ID_U is the identity map on $F(U)$.

An element $\sigma \in F(U)$ is called a SECTION of F on U , and an \mathcal{R} -linear map ρ_{VU} is called a RESTRICTION map. We often write $\sigma|_V$ instead of $\rho_{VU}(\sigma)$, and call it the restriction of σ to V .

Definition 2 (Sheaf). A presheaf F on X is called a SHEAF if it satisfies the following two conditions:

1. For any open set $U \subset X$, any open covering $U = \cup_{i \in I} U_i$, and any section $\sigma \in F(U)$, $\sigma|_{U_i} = 0$ for all $i \in I$ implies $\sigma = 0$.
2. For any open set $U \subset X$, any open covering $U = \cup_{i \in I} U_i$, any family $\sigma_i \in F(U_i)$ satisfying $\sigma_i|_{U_i \cap U_j} = \sigma_j|_{U_i \cap U_j}$ for all pairs (i, j) , there exists $\sigma \in F(U)$ such that $\sigma|_{U_i} = \sigma_i$ for all $i \in I$.

Each \mathcal{R} -module $F(U)$ is best regarded as “local data” on U . From the conditions in Definition 2, a sheaf F allows one to glue a set of local data together into global data uniquely. This hints at its utility.

2.2 Network Coding Sheaves

We recall the problem setting of network coding [2, 11, 15]. Let k be an \mathcal{R} -module, or simply a (finite) field. Let $G = (\mathcal{V}, \mathcal{E})$ be a directed graph (not necessarily acyclic), where \mathcal{V} and \mathcal{E} are finite sets of nodes and directed edges, respectively. A directed edge $e \in \mathcal{E}$ from $v \in \mathcal{V}$ to $w \in \mathcal{V}$ is denoted by $e = |vw|$ ($\text{HEAD}(e) := w$, $\text{TAIL}(e) := v$). All graphs in this paper are viewed as topological spaces with the usual locally Euclidean cellular topology.

We assume that there exists a subset $S = \{s_1, \dots, s_\alpha\} \subset \mathcal{V}$ of nodes called SOURCES which transmit, as information, elements in $k^{n_{s_i}}$, $n_{s_i} \in \mathbb{N}$, for each $s_i \in S$. We also assume that there exists a subset $R = \{r_1, \dots, r_\beta\} \subset \mathcal{V}$ of nodes called RECEIVERS. Each receiver requires information from some sources and this assignment is determined by $\mathcal{S} : R \rightarrow 2^S$ in the sense that a receiver r_i requires all transmitted information from $\mathcal{S}(r_i) \in 2^S$.

Let $\text{cap} : \mathcal{E} \rightarrow \mathbb{N}$ be a CAPACITY function which assigns for each edge $e \in \mathcal{E}$ its information capacity $\text{cap}(e)$. The set of the incoming (outgoing, resp.) edges in the sense of edge directions at a node $v \in \mathcal{V}$ is denoted by $\text{In}(v)$ ($\text{Out}(v)$, resp.). A LOCAL CODING MAP, ϕ_{vw} determines a data assignment of the incoming data at v into an outgoing edge e with $\text{HEAD}(e) = w$ given by

$$\phi_{vw} : k^{n_v} \oplus k^{l_v} \rightarrow k^{\text{cap}(e)}, \text{ where } l_v = \sum_{e \in \text{In}(v)} \text{cap}(e),$$

where it is assumed that $n_v = 0$ for $v \in \mathcal{V} \setminus S$. In particular, a local coding map $\phi_{s_i r_j}$ from a receiver r_j to a source $s_i \in \mathcal{S}(r_j)$ corresponds to the DECODING MAP. The local coding map at v is the linear transformation Φ_v sending incoming data to outgoing data; its row decomposition is precisely the ϕ_{vw} above. Denote the set of all local coding maps by $\Phi = \{\phi_{vw}\}$.

In order to express decodable information flows on a network as a (co)cycle, we complete the graph $G = (\mathcal{V}, \mathcal{E})$ to $X = (\mathcal{V}, \tilde{\mathcal{E}})$, where $\tilde{\mathcal{E}}$ is given by adding edges $e = |r_j s_i|$ in \mathcal{E} from each receiver r_j to all of its requesting sources $s_i \in \mathcal{S}(r_j)$ with $\text{cap}(e = |r_j s_i|) = n_{s_i}$. To remove ambiguity, we denote the set of incoming edges at $v \in \mathcal{V}$ in \mathcal{E} or $\tilde{\mathcal{E}}$ by $\text{In}(v; \mathcal{E})$ or $\text{In}(v; \tilde{\mathcal{E}})$, respectively, with $\text{Out}(v; \mathcal{E})$ and $\text{Out}(v; \tilde{\mathcal{E}})$ similarly defined. This extension enables one to compare decoded information at each receiver r_j with transmitted information from $s_i \in \mathcal{S}(r_j)$ as the gluing condition of the network coding sheaf on the added edge $e = |r_j s_i|$.

We define the network coding (NC) sheaf F associated to (X, Φ) locally, as follows:

Definition 3 (Local Sections).

1. For a connected open set U contained in an edge $e \in \tilde{\mathcal{E}}$, $F(U) := k^{\text{cap}(e)}$.
2. For a connected open set U which only contains one node $v \in \mathcal{V}$, $F(U) := k^{n_v} \oplus k^{l_v}$, where $l_v = \sum_{e \in \text{In}(v; \tilde{\mathcal{E}})} \text{cap}(e)$.

Definition 4 (Local Restriction Maps).

1. For connected open sets $V \subset U \subset e$ for some edge e , $\rho_{VU} := \text{ID} : F(U) \rightarrow F(V)$.
2. For connected open sets $V \subset U$, where U contains only one node v and V is located in $e \in \text{In}(v; \tilde{\mathcal{E}})$, $\rho_{VU} : F(U) \rightarrow F(V)$ is given by the projection map induced by the product structure in Definition 3.
3. For connected open sets $W \subset U$, where U contains only one node v and W is located in $e \in \text{Out}(v; \tilde{\mathcal{E}})$, $\rho_{WU} := \phi_{vw} : F(U) \rightarrow F(W)$, where $w = \text{HEAD}(e)$.

From these local definitions of sections and restriction maps, the network coding sheaf is defined by the SHEAFIFICATION [1], a standard process to construct sections and restriction maps via the gluing condition in Definition 2, specifically:

Definition 5 (NC Sheaf). For $U \subset X$, $F(U)$ is defined to be the set of all equivalent classes $\sigma = [(\sigma_i, U_i)_{i \in I}]$, where a representative $(\sigma_i, U_i)_I$ with a covering $U = \cup_I U_i$ is given by a family of sections $\sigma_i \in F(U_i)$ satisfying $\sigma_i|_{U_i \cap U_j} = \sigma_j|_{U_i \cap U_j}$, and the equivalent relation \sim is defined by

$$(\sigma_i, U_i)_I \sim (\tau_j, V_j)_J \Leftrightarrow \sigma_i|_{U_i \cap V_j} = \tau_j|_{U_i \cap V_j}, \text{ for } i \in I, j \in J.$$

The restriction map $\rho_{VU} : F(U) \rightarrow F(V)$ is induced by local restriction maps on a representative (independent of the choice of a representative). The sheaf F obtained by the sheafification process is called the network coding sheaf of (X, Φ) .

2.3 Sheaf Cohomology

Cohomology is a basic invariant of topological spaces which captures global features of the space by means of homological algebra [7]. In like manner, for a sheaf F taking values in \mathcal{R} -modules, the global structure of the sheaf data on X is characterized by its SHEAF COHOMOLOGY, $H^\bullet(X; F)$, a graded \mathcal{R} -module. General sheaf cohomology is too involved to describe in this short article [8, 10]; we resort to a more limited (and, fortunately, equivalent and computable) variant, Čech cohomology.

To define Čech cohomology of F , choose the open covering $X = (\cup_{v \in \mathcal{V}} U_v) \cup (\cup_{e \in \tilde{\mathcal{E}}} U_e)$ by using open stars U_v and U_e for each $v \in \mathcal{V}$ and $e \in \tilde{\mathcal{E}}$. Here, an open star U_v for a node $v \in \mathcal{V}$ is the maximal connected open set containing only one node v , and an open star U_e for an edge $e \in \tilde{\mathcal{E}}$ is the maximal open set contained in the edge e . Define the Čech complex $0 \rightarrow C^0(X; F) \xrightarrow{\partial} C^1(X; F) \rightarrow 0$ as:

$$C^0(X; F) = \prod_{v \in \mathcal{V}} F(U_v) \quad ; \quad C^1(X; F) = \prod_{e \in \tilde{\mathcal{E}}} F(U_e), \quad (2.1)$$

where the boundary map $\partial = (\partial_e)_{e \in \tilde{\mathcal{E}}}$ is defined for each product element $F(U_e)$ of $C^1(X; F)$ with $e = |vw|$ by

$$\begin{aligned} \partial_e : F(U_v) \times F(U_w) &\rightarrow F(U_e), \\ \partial_e(\sigma_v, \sigma_w) &= \rho_{U_e U_v}(\sigma_v) - \rho_{U_e U_w}(\sigma_w). \end{aligned} \quad (2.2)$$

Definition 6 (Sheaf Cohomology). The i -th sheaf cohomology $H^i(X; F)$ is defined by $H^i(X; F) := H^i(C^\bullet)$, i.e.,

$$H^0(X; F) := \ker(\partial), \quad (2.3)$$

$$H^1(X; F) := C^1(X; F)/\text{im}(\partial). \quad (2.4)$$

For an open set $A \hookrightarrow X$, a sheaf F on X induces a sheaf on A called the INVERSE IMAGE ι^*F . It is defined by $\iota^*F(U) := F(U)$ for an open set $U \subset A$, and the restriction maps are induced by the original ones of the sheaf F . Then, by constructing an open covering for A from $X = (\cup_{v \in \mathcal{V}} U_v) \cup (\cup_{e \in \tilde{\mathcal{E}}} U_e)$, one can define the sheaf cohomology $H^\bullet(A; \iota^*F)$ on A as the cohomology

of the Čech complex $C^\bullet(A; \iota^*F)$. We will often use the notations $H^\bullet(A; F) = H^\bullet(A; \iota^*F)$ and $C^\bullet(A; F) = C^\bullet(A; \iota^*F)$.

The relative sheaf cohomology $H^\bullet(X, A; F)$ with respect to $A \subset X$ open is defined as follows. Any A open defines a surjective chain map $p^\bullet : C^\bullet(X; F) \rightarrow C^\bullet(A; F)$. The relative chain complex is defined as the subcomplex $C^\bullet(X, A; F) := \ker(p^\bullet)$. The relative sheaf cohomology is defined as $H^\bullet(X, A; F) := H^\bullet(C^\bullet(X, A; F))$.

2.4 Computation

It should be noted that computations for NC sheaf cohomology require only module operations, thanks to the use of Čech cohomology. In particular, from the definition of the sheaf cohomology (2.3),(2.4), it suffices to check the kernel and the cokernel of the boundary map $\partial : C^0(X; F) \rightarrow C^1(X; F)$. These calculations are performed by means of Smith normal forms. We refer to [9] for details of computations of Smith normal forms including fast algorithms and reduction pre-processing.

2.5 Information-Theoretic Content of H^\bullet

The cohomologies of NC sheaves provide a concise global collation of algebraic, topological, and information-theoretic content. We begin with an interpretation of $H^0(X; F)$ for a NC sheaf F . From (2.3), it suffices to examine the kernel of the boundary map $\partial : C^0(X; F) \rightarrow C^1(X; F)$.

We recall the definition of an information flow on a network G with coding Φ . An INFORMATION FLOW ψ for a family of transmitted data $z = (z_{s_1}, \dots, z_{s_\alpha})$, $z_{s_i} \in k^{n_{s_i}}$, $s_i \in S$, is defined by an assignment $\psi(e) \in k^{\text{cap}(e)}$ for each edge $e \in \mathcal{E}$ satisfying the FLOW CONDITIONS: the data in ψ are related by local coding maps Φ_v at all vertices v . More specifically, for $e = |vw|$ and $e_i \in \text{In}(v; \mathcal{E})$ ($i = 1, \dots, K$),

1. $\phi_{vw}(\psi(e_i)_\uparrow^K) = \psi(e)$ for $v \notin S \cup R$,
2. $\phi_{vw}(z_{s_i}, \psi(e_i)_\uparrow^K) = \psi(e)$ for $v = s_i \in S$,
3. $\phi_{vw}(\psi(e_i)_\uparrow^K) = z_{s_i}$ for $v = r_j \in R$, $w = s_i \in S(r_j)$,
4. $\phi_{vw}(\psi(e)_\uparrow^K) = \psi(e)$ for $v = r_j \in R$, $w \notin S(r_j)$,

and so on. The other cases are similarly derived by taking proper domain and target spaces of local coding maps.

We recall that the boundary map $\partial : C^0(X; F) \rightarrow C^1(X; F)$ is determined by a family of maps ∂_e ($e \in \tilde{\mathcal{E}}$) by (2.2). Hence, $\sigma = (\sigma_v)_{v \in V} \in C^0(X; F) \in \ker(\partial)$ if and only if $\partial_e(\sigma_v, \sigma_w) = 0$ for all $e = |vw| \in \tilde{\mathcal{E}}$. The restriction map $\rho_{U_e U_v}$ from the tail node $v = \text{TAIL}(e)$ is determined by the local coding map ϕ_{vw} , and the restriction map $\rho_{U_e U_w}$ from the head node $w = \text{HEAD}(e)$ is determined by the projection $F(U_w) \rightarrow F(U_e)$.

Theorem 7 (Information Theoretic Content of $H^0(X; F)$). *For a NC sheaf F of (X, Φ) , elements of the sheaf cohomology $H^0(X; F)$ are in bijective correspondence with compatible information flows on the network.*

Proof. Construct an assignment $\psi(e) \in k^{\text{cap}(e)}$ for each $e = |vw| \in \mathcal{E}$ by $\psi(e) = \rho_{U_e U_w}(\sigma_w)$ (i.e., projecting $\sigma_w \in F(U_w)$ into $F(U_e)$). Let $z = (z_{s_1}, \dots, z_{s_\alpha})$, $z_{s_i} \in k^{n_{s_i}}$, be the product element

in $\sigma \in C^0(X; F)$ corresponding to the transmitted data from S . A simple direct calculation reveals that the kernel condition on ∂_e expressed as $\rho_{u_e u_v}(\sigma_v) = \rho_{u_e u_w}(\sigma_w)$ is equivalent to the above flow condition for the assignment ψ on the edge e . We therefore obtain the equivalent expression of an information flow for a family of transmitted data z by means of $\sigma \in \ker(\partial)$. \square

This theorem makes it possible to apply homological-algebraic tools for sheaf cohomology to network coding problems.

3 Applications

3.1 Relative Cohomology and Maxflow Bounds

Recall the definition of relative NC sheaf cohomology $H^\bullet(X, A; F)$ for an open set $A \subset X$: the relative chain complex $C^\bullet(X, A; F)$ is derived as a subcomplex of $C^\bullet(X; F)$ which is mapped to 0 by the surjective chain map $p^\bullet : C^\bullet(X; F) \rightarrow C^\bullet(A; F)$. Hence, we have a short exact sequence of chain complexes

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C^0(X, A; F) & \xrightarrow{i^0} & C^0(X; F) & \xrightarrow{p^0} & C^0(A; F) & \longrightarrow & 0 \\ & & \downarrow \partial & \circ & \downarrow \partial & \circ & \downarrow \partial & & \\ 0 & \longrightarrow & C^1(X, A; F) & \xrightarrow{i^1} & C^1(X; F) & \xrightarrow{p^1} & C^1(A; F) & \longrightarrow & 0 \end{array} \quad (3.1)$$

meaning that the above diagram is commutative and $\ker(i^k) = 0$, $\text{im}(i^k) = \ker(p^k)$, $\text{im}(p^k) = C^k(A; F)$ for $k = 0, 1$. The boundary maps in (3.1) are all induced by the original one $\partial : C^0(X; F) \rightarrow C^1(X; F)$. One of the important techniques in homological algebra is the LONG EXACT SEQUENCE induced by a short exact sequence. For (3.1), the induced long exact sequence is:

$$0 \rightarrow H^0(X, A; F) \xrightarrow{i^0} H^0(X; F) \xrightarrow{p^0} H^0(A; F) \xrightarrow{\delta^0} H^1(X, A; F) \xrightarrow{i^1} \dots \quad (3.2)$$

where the maps i^\bullet and p^\bullet in (3.2) are induced by those in the short exact sequence (3.1). The map δ^0 is called the CONNECTING HOMOMORPHISM and is given by $\delta^0(\sigma_A) := \partial(\sigma_X)$, where $p^0(\sigma_X) = \sigma_A$; furthermore, $\partial(\sigma_X) \in C^1(X; F)$ is identified with the element in $C^1(X, A; F)$, since $p^1 \partial(\sigma_X) = \partial p^0(\sigma_X) = \partial \sigma_A = 0$ leads to $\partial(\sigma_X) \in C^1(X, A; F)$.

An elementary application of the long exact sequence (3.2) yields information flow capacity bounds. This is most transparent in the single-source scenario, $S = \{s\}$. Consider an open set $A \subset X - s$ which does not include the source node, but does include some receiver $r_j \in R$. Then, the union of incoming edges into A (those contained neither entirely in A or its complement) define a CUT C_A between s and r_j ; and any cut may be realized as coming from some open subset A as above. Recall that the CAPACITY of a cut C_A is the sum of the edge capacities over all edges in C_A .

Lemma 8. For $A \subset X - s$ containing a receiver and F any NC sheaf:

1. $\text{cap}(C_A) = \dim H^0(A; F)$;

$$2. H^0(X, A; F) = 0.$$

Corollary 9. *The maxflow is bounded below by the mincut.*

Proof. From Lemma 8, $H^0(X, A; F) = 0$ in (3.2) and hence $p^0 : H^0(X; F) \rightarrow H^0(A; F)$ is injective, so that $\dim H^0(X; F) \leq \dim H^0(A; F)$ for any network coding sheaf F and any cut-inducing A . By Theorem 7, $\dim H^0(X; F)$ equals the flow capacity of the network coding for F . Maximizing and minimizing over all NC sheaves F yields,

$$\overbrace{\max_F \dim H^0(X; F)}^{\text{maxflow}} \leq \overbrace{\min_A \dim H^0(A; F)}^{\text{mincut}}.$$

□

This is, of course, superfluous and more easily proved with less cumbersome tools. A cohomological proof, however, may apply to situations which are currently unknown or difficult to analyze.

3.2 Extensions and the Connecting Morphism

The long exact sequence of a pair (X, A) (3.2), examined in more detail and for general $A \subset X$ open, reveals more structure in the multi-source setting. Of relevance is the situation where one fixes a local information flow σ_A on A and studies the EXTENSION PROBLEM: does σ_A extend globally to a network flow respecting the coding and capacity constraints?

Proposition 10 (Global Extendability). *A local information flow $\sigma_A \in H^0(A; F)$ is globally extendable if and only if it lies in the kernel of the connecting homomorphism: $\delta^0(\sigma_A) = 0$.*

Proof. Since the information flows on A and X are represented by elements in $H^0(A; F)$ and $H^0(X; F)$, respectively, it suffices to find $\sigma \in H^0(X; F)$ such that $p^0(\sigma) = \sigma_A$. The exactness of the sequence (3.2) at $H^0(A; F)$ is expressed as $\text{im}(p^0) = \ker(\delta^0)$. □

3.3 Excision and Network Robustness

Let $A \subset G$ be an open set and $Z = X \setminus A$ be its complementary closed set. For a section $\sigma \in F(U)$, the support of σ , denoted $|\sigma|$, is defined as

$$|\sigma| := \{x \in U \mid s|_V \neq 0 \text{ for any neighborhood } V \subset U \text{ of } x\}.$$

Consider also the subspace of $F(U)$ for each open set $U \subset X$ given by

$$F_Z(U) := \{\sigma \in F(U) \mid |\sigma| \subset Z\}.$$

Then by replacing $F(U_v)$, $F(U_e)$ in (2.1) with $F_Z(U_v)$, $F_Z(U_e)$, respectively, the local cohomology with support Z , denoted $H_Z^\bullet(X; F)$, is defined in the same way. The local cohomology $H_Z^\bullet(X; F)$ is expressed in an exact sequence as follows [8, II.9.2]:

$$0 \rightarrow H_Z^0(X; F) \xrightarrow{i^0} H^0(X; F) \xrightarrow{p^0} H^0(A; F) \xrightarrow{\delta^0} H_Z^1(X; F) \xrightarrow{i^1} \dots \quad (3.3)$$

where i^* is induced by the inclusion map $i_U : F_Z(U) \rightarrow F(U)$.

Suppose that a network experiences a failure on a sub-network $A \subset G$. The ROBUSTNESS PROBLEM for information flows asks under which the global information flow $\sigma \in H^0(X; F)$ persists on Z with the removal of A .

Proposition 11 (Network Robustness). *Let $A \subset G$ be an open set and $Z = X \setminus A$ be the complementary closed set. Then $H_Z^0(X; F)$ represents the global information flow on the failure network $G \setminus A$. Moreover, the network coding of F is robust to this failure if and only if $p^0 = 0$.*

Proof. From the definition of the local cohomology $H_Z^0(X; F)$, a section $\sigma_Z \in H_Z^0(X; F)$ represents an information flow on the network G such that $|\sigma_Z| \subset Z$, i.e., for any point $x \in A$, there exists an open set $U \subset A$ such that $\sigma_Z|_U = 0$. Hence, it follows that $H_Z^0(X; F)$ represents the information flows which are not changed by the removal of A . The exact sequence (3.3) completes the argument. \square

Remark 12. By the Five Lemma [4], we have the isomorphism $H^k(X, A; F) \simeq H_Z^k(X; F)$. Hence the above argument can be explained by using only relative cohomology. On the other hand, the long exact sequence (3.3) is one of the examples showing the EXCISION PROPERTY. There are several versions of long exact sequences related to excision property [8, II.9], each of which can be used to analyze local information flows as above.

3.4 Mayer-Vietoris and Data Merging

In this subsection, we study a data merging problem via homological algebra. Let U, V be open sets in X such that $X = U \cup V$. The problem of DATA MERGING — whether local information flows on U and V can be merged to a global information flow on X — is amenable to an exact sequence interpretation.

Proposition 13. (Data Merging). *Let U and V be open sets in X and $\sigma_U \in H^0(U; F)$ and $\sigma_V \in H^0(V; F)$ be local information flows on U and V , respectively. Then these two local information flows can be merged into a global information flow on X if and only if the local information flows agree (are cohomologous) on $U \cap V$.*

Proof. The Mayer-Vietoris long exact sequence [8, II.5.10] states that:

$$\begin{aligned} 0 \rightarrow H^0(X; F) \xrightarrow{f^0} H^0(U; F) \oplus H^0(V; F) \xrightarrow{g^0} H^0(U \cap V; F) \xrightarrow{\delta^0} \\ \xrightarrow{\delta^0} H^1(X; F) \xrightarrow{f^1} H^1(U; F) \oplus H^1(V; F) \xrightarrow{g^1} \dots, \end{aligned} \quad (3.4)$$

where f^* and g^* are defined by the sum and difference of restriction maps, respectively. Let $\sigma_U \in H^0(U; F)$ and $\sigma_V \in H^0(V; F)$ be local information flows on U and V , respectively. By definition, the existence of a global information flow $\sigma \in H^0(X; F)$ which merges σ_U and σ_V implies that $f^0(\sigma) = (\sigma_U, \sigma_V)$. By exactness, this is equivalent to the pair of local information flows being in the kernel of g_0 , i.e., their difference on $U \cap V$ is null. \square

4 Conclusion

This paper marks the introduction of sheaf-theoretical tools to network coding. Sheaf cohomology and associated exact sequences are shown to be useful in several network coding contexts.

Due to space limitations, the full scope of our proposed program and full details of proofs are not presented.

We anticipate one important application of sheaf theory in our future work to be a characterization of maxflows on general multi-source network codings. One of the possibilities to attack this problem is to give another sheaf theoretic proof of the maxflow characterization on a single source scenario, which will be potentially applicable to multi-source one. As is seen in the single source case, the maxflow characterization given by the mincut can be regarded as a flow-cut duality. From this viewpoint of duality, a derived categorical formulation of network coding sheaves may provide us with some useful characterizations via Poincaré - Verdier duality and Morse theory.

It should be also mentioned that we can differently formulate network coding problems without the extension of the graph $G \rightarrow X$. Even in this formulation, information flows on a network can be treated by using NC sheaf cohomologies [6]. Besides cohomology, other basic operations on sheaves (e.g., f_* , f^* , \otimes , $\mathcal{H}om(\bullet, \bullet)$) can be defined [4, 8, 10] and are useful for constructing new NC sheaves or to investigate relationships between different NC sheaves and their cohomologies. Because of the generality of the sheaf cohomological tools presented here, extensions to higher dimensional base spaces are straightforward. We propose this to be useful in situations with spacial expanse (wireless broadcast) or with time dependence (using the time axis \mathbb{R} in the base space).

Acknowledgment

The authors express their sincere gratitude to S. Hara and M. Robinson for valuable discussions on this paper. This work is partially supported by AFOSR # FA9550-09-1-0643 and DARPA # HR0011-07-1-0002 for the first author, and by JST PRESTO program for the second author.

References

- [1] G. Bredon, *Sheaf Theory*, Springer, 1997.
- [2] C. Fragouli and E. Soljanin. *Network Coding Fundamentals*. *Foundations and Trends in Networking*, Vol. 2, Issue 1:1-133, 2007.
- [3] C. Fragouli and E. Soljanin. *Network Coding Applications*. *Foundations and Trends in Networking*, Vol. 2, Issue 2:135-269, 2007.
- [4] S. I. Gelfand and Y. I. Manin, *Methods of Homological Algebra*, 2nd Edition, Springer, 2003.
- [5] J. A. Goguen, *Sheaf Semantics for Concurrent Interacting Objects*, *Math. Structures Comput. Sci.* 2 (1992), no. 2, 159–191.
- [6] R. Ghrist and Y. Hiraoka, *A Sheaf Cohomological Approach to Network Codings*, in preparation.
- [7] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [8] B. Iversen, *Cohomology of Sheaves*, Springer, 1986.

- [9] T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational Homology*, Springer, 2004.
- [10] M. Kashiwara and P. Schapira, *Sheaves on Manifolds*, Springer, 1990.
- [11] R. Koetter and M. Médard, An Algebraic Approach to Network Coding, *IEEE Trans. on Networking*, Vol. 11 (2003), 782–795.
- [12] M. Robinson, Inverse problems in geometric graphs using internal measurements, [arXiv:1008.2933v1](https://arxiv.org/abs/1008.2933v1), August 2010.
- [13] M. Robinson, Asynchronous logic circuits and sheaf obstructions, [arXiv:1008.2729v1](https://arxiv.org/abs/1008.2729v1), August 2010.
- [14] E. Şendroiş, Sheaf Tools for Computation, *Appl. Math. Comput.* 184 (2007), no. 1, 131–141.
- [15] R. W. Yeung, *A First Course in Information Theory*, Springer, 2002.