An Algebraic Approach to Network Coding

July 30–31, 2009

Outline

Outline

An Example

Recall that in a linear network the random process on an edge is given by

$$Y(e) = \sum_{k=1}^{\mu(v)} \alpha_{k,e} X(v,k) + \sum_{e':\text{head}(e')=\text{tail}(e)} \beta_{e',e} Y(e')$$

and that the output processes are given by

$$Z(v,k) = \sum_{e':\text{head}(e')=v} \varepsilon_{e',k} Y(e').$$

Connection Types

The various types of connections $\mathcal{C}$:

- **Point-to-Point:** $\mathcal{C} = \{(v, u, \mathcal{X}(v, u))\}$

Connection Types

The various types of connections $\mathcal{C}$:

- **Point-to-Point:** $\mathcal{C} = \{\big(v, u, \mathcal{X}(v, u)\big)\}$
- **Multicast:** $\mathcal{C} = \{\big(v, u_j, \mathcal{X}(v, u_j)\big) \mid j = 1, \ldots, K\}$

Connection Types

The various types of connections $\mathcal{C}$:

- **Point-to-Point:** $\mathcal{C} = \{(v, u, \mathcal{X}(v, u))\}$
- **Multicast:** $\mathcal{C} = \{(v, u_j, \mathcal{X}(v, u_j)) \mid j = 1, \ldots, K\}$
- **General:** $\mathcal{C} = \{(v_i, u_j, \mathcal{X}(v_i, u_j)) \mid i = 1, \ldots, N \text{ and } j = 1, \ldots, K\}$

Outline

Capacities and Cuts

With each edge $e \in E$, we associated a non-negative number $C(e)$, called the capacity of e. In our setup, we will assume that $C(e) = C(e')$ for all $e, e' \in E$.

Capacities and Cuts

With each edge $e \in E$, we associated a non-negative number $C(e)$, called the capacity of e. In our setup, we will assume that $C(e) = C(e')$ for all $e, e' \in E$.

Definition

A cut between vertices v and v' in $\mathcal{G} = (V, E)$ is a partition of V into two classes S and $S^c = V \setminus S$ such that $v \in S$ and $v' \in S^c$. The value of a cut is defined as

$$V(S) = \sum_{e \in [S, S^c]} C(e)$$

where $C(e)$ denotes the capacity of a link and $[S, S^c] := \{e \in E \mid \text{tail}(e) \in S \text{ and } \text{head}(e) \in S^c\}$.

Min-Cut Max-Flow

Theorem (Min-Cut Max-Flow)

Let a network with a single connection $c = (v, v', \mathcal{X}(v, v'))$ *be given. Then the network problem is solvable if and only if the rate* $R(c)$ *of the connection is less than or equal to the minimum value of all cuts between* v *and* v'*.*

Zeros Abound

Lemma

Let $\mathbb{F}[X_1, \ldots, X_n]$ *be the ring of polynomials over an infinite field* $\mathbb{F}$ *in variables* $X_1, \ldots, X_n$*. For any nonzero element* $f \in \mathbb{F}[X_1, \ldots, X_n]$*, there exists an infinite set of* n*-tuples* $(x_1, \ldots, x_n) \in \mathbb{F}^n$ *such that* $f(x_1, \ldots, x_n) \neq 0$*.*

Zeros Abound

Lemma

Let $\mathbb{F}[X_1, \ldots, X_n]$ be the ring of polynomials over an infinite field $\mathbb{F}$ in variables $X_1, \ldots, X_n$. For any nonzero element $f \in \mathbb{F}[X_1, \ldots, X_n]$, there exists an infinite set of n-tuples $(x_1, \ldots, x_n) \in \mathbb{F}^n$ such that $f(x_1, \ldots, x_n) \neq 0$.

This lemma will be useful because we will be considering the algebraic closure $\bar{\mathbb{F}}$ of $\mathbb{F}_2$, which turns out to be

$$\bar{\mathbb{F}} = \bigcup_{m \in \mathbb{N}} \mathbb{F}_{2^m}.$$

A Nice Equivalence

Theorem

Let a linear network be given with source node v, sink node v', and a desired connection $c = (v, v', \mathcal{X}(v, v'))$ of rate $R(c)$. The following three statements are equivalent:

1. *A point-to-point connection $c = (v, v', \mathcal{X}(v, v'))$ is possible.*
2. *The Min-Cut Max-Flow bound is satisfied between v and v' for a rate $R(c)$.*
3. *The determinant of the $R(c) \times R(c)$ transfer matrix M is nonzero over the ring $\mathbb{F}_2[\ldots, \alpha_{l,e}, \ldots, \beta_{e',e}, \ldots, \varepsilon_{e',j}, \ldots]$.*

A Nice Equivalence

Theorem

Let a linear network be given with source node v, sink node v', and a desired connection $c = (v, v', \mathcal{X}(v, v'))$ of rate $R(c)$. The following three statements are equivalent:

1. *A point-to-point connection $c = (v, v', \mathcal{X}(v, v'))$ is possible.*
2. *The Min-Cut Max-Flow bound is satisfied between v and v' for a rate $R(c)$.*
3. *The determinant of the $R(c) \times R(c)$ transfer matrix M is nonzero over the ring $\bar{\mathbb{F}}[\dots, \alpha_{l,e}, \dots, \beta_{e',e}, \dots, \varepsilon_{e',j}, \dots]$.*

In the proof that's given, we seem to need an infinite field for the implication $3 \implies 1$.

Outline

Line Graphs

Definition

We define the directed labeled line graph of $\mathcal{G} = (V, E)$ as $\mathfrak{G} = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V} = E$ and edge set

$$\mathcal{E} = \{(e, e') \in E \times E \mid \text{head}(e) = \text{tail}(e')\}.$$

An edge $(e, e') \in \mathcal{E}$ will be labeled with $\beta_{e',e}$.

Adjacency Matrices

Given $\mathcal{G} = (V, E)$, define the adjacency matrix F of the line graph $\mathfrak{G}$ to be the $|E| \times |E|$ matrix with entries

$$F_{i,j} = \begin{cases} \beta_{e_i,e_j}, & \text{head}(e_i) = \text{tail}(e_j); \\ 0, & \text{otherwise.} \end{cases}$$

Adjacency Matrices

Given $\mathcal{G} = (V, E)$, define the adjacency matrix F of the line graph $\mathfrak{G}$ to be the $|E| \times |E|$ matrix with entries

$$F_{i,j} = \begin{cases} \beta_{e_i,e_j}, & \text{head}(e_i) = \text{tail}(e_j); \\ 0, & \text{otherwise.} \end{cases}$$

Lemma

Let F be the adjacency matrix of the labeled line graph of a cycle-free network $\mathcal{G}$. The matrix $I - F$ has a polynomial inverse with coefficients in $\mathbb{F}_2[\ldots, \beta_{e',e}, \ldots]$.

Input and Output

We consider the vectors $\underline{x} = (x_1, \ldots, x_\mu)$ and $\underline{z} = (z_1, \ldots, z_\nu)$ of all input and output processes, respectively. So, for any i, $x_i = X(v, l)$ for some $v \in V$ and some $l \in \mathbb{N}$. Similarly, $z_j = Z(v, l)$.

Input and Output

We consider the vectors $\underline{x} = (x_1, \ldots, x_\mu)$ and $\underline{z} = (z_1, \ldots, z_\nu)$ of all input and output processes, respectively. So, for any i, $x_i = X(v, l)$ for some $v \in V$ and some $l \in \mathbb{N}$. Similarly, $z_j = Z(v, l)$.

We now define the $\mu \times |E|$ matrix A to have entries

$$A_{i,j} = \begin{cases} \alpha_{l,e_j}, & x_i = X(\mathrm{tail}(e_j), l); \\ 0, & \text{otherwise}; \end{cases}$$

and the $\nu \times |E|$ matrix B to have entries

$$B_{i,j} = \begin{cases} \varepsilon_{e_j,l}, & z_i = Z(\mathrm{head}(e_j), l); \\ 0, & \text{otherwise}. \end{cases}$$

The Form of a Transfer Matrix

Theorem

Let a network be given with matrices A, B, and F. The transfer matrix of the network is given as

$$M = A(I - F)^{-1}B^T$$

where I is the $|E| \times |E|$ identity matrix.

Outline

Multicast

Theorem

Let a delay-free network $\mathcal{G}$ *and a set of desired connections* $\mathcal{C} = \{(v, u_i, \mathcal{X}(v)) \mid i = 1, \dots, N\}$ *be given. The network problem* $(\mathcal{G}, \mathcal{C})$ *is solvable if and only if the Min-Cut Max-Flow bound is satisfied for all connections in* $\mathcal{C}$*.*

Multicast

Theorem

Let a delay-free network $\mathcal{G}$ and a set of desired connections $\mathcal{C} = \{(v, u_i, \mathcal{X}(v)) \mid i = 1, \ldots, N\}$ be given. The network problem $(\mathcal{G}, \mathcal{C})$ is solvable if and only if the Min-Cut Max-Flow bound is satisfied for all connections in $\mathcal{C}$.

An important part of this theorem is that all sink nodes get the same information, in this case $\mathcal{X}(v)$.

A Bound on the Size of the Base Field

For convenience, we will let $\underline{\xi} = \{\xi_1, \ldots, \xi_n\}$ be the collection of all variables of the forms α_{l,e_j}, $\beta_{e',e}$, and $\varepsilon_{e_j,l}$.

A Bound on the Size of the Base Field

For convenience, we will let $\underline{\xi} = \{\xi_1, \ldots, \xi_n\}$ be the collection of all variables of the forms α_{l,e_j}, $\beta_{e',e}$, and $\varepsilon_{e_j,l}$.

Theorem

Let a delay-free communication network $\mathcal{G}$ and a solvable multicast network problem be given with one source and N rreceivers. Let F be the product of the determinants of the transfer matrices for the individual connections and let δ be the maximal degree of F with respect to any variable ξ_i. There exists a solution to the multicast network problem in $\mathbb{F}_{2^i}$, where i is the smallest number such that $2^i > \delta$. Moreover, there is a simple greedy algorithm that finds such a solution.

Corollary

Let a delay-free communication network $\mathcal{G}$ and a solvable multicast network problem be given with one source and N recievers. Let R be the rate at which the source generates information. There exists a solution to the network coding problem in a finite field $\mathbb{F}_{2^m}$ with $m \leq \lceil \log_2(NR+1) \rceil$.

Outline

General Networks

Theorem (Generalized Min-Cut Max-Flow Condition)

Let an acyclic delay-freee linear network problem $(\mathcal{G}, \mathcal{C})$. *be given, and let* $M = \{M_{i,j}\}$ *be the corresponding tranfer matrix relating the set of input nodes to the set of output nodes. The network problem is solvable if and only if there exists an assignment of numbers to variables* $\underline{\xi}$ *such that*

1. $M_{i,j} = 0$ *for all pairs* (v_i, v_j) *of vertices such that* $(v_i, v_j, \mathcal{X}(v_i, v_j)) \notin \mathcal{C}$.
2. *if* $\mathcal{C}$ *contains the connections* $(v_{i_k}, v_j, \mathcal{X}(v_i, v_j))$ *for* $k = 1, \ldots, l$, *then the submatrix* $[M^T_{i_1,j}, \ldots, M^T_{i_l,j}]$ *is a nonsingular* $\nu(v_j) \times \nu(v_j)$ *matrix.*

A Little Algebraic Geometry

Let $f_1(\underline{\xi}), \ldots, f_K(\underline{\xi})$ denote all of the entries in M that have to evaluate to zero in order to satisfy the first condition of the Generalize Min-Cut Max-Flow theorem. We consider the ideal generated by $f_1(\underline{\xi}), \ldots, f_K(\underline{\xi})$ and denote this ideal by $I(f_1, \ldots, f_K)$.

A Little Algebraic Geometry

Let $f_1(\underline{\xi}), \ldots, f_K(\underline{\xi})$ denote all of the entries in M that have to evaluate to zero in order to satisfy the first condition of the Generalize Min-Cut Max-Flow theorem. We consider the ideal generated by $f_1(\underline{\xi}), \ldots, f_K(\underline{\xi})$ and denote this ideal by $I(f_1, \ldots, f_K)$.

Let $g_1(\underline{\xi}), \ldots, g_L(\underline{\xi})$ denote the determinants of the $\nu(v_j) \times \nu(v_j)$ matrices that have to be nonzero. We introduce a new variable ξ_0 and consider the funtion $\xi_0 \prod_{i=1^L} g_i(\underline{\xi}) - 1$. We call the ideal $I(f_1(\underline{\xi}), \ldots, f_K(\underline{\xi}), \xi_0 \prod_{i=1^L} g_i(\underline{\xi}) - 1)$ the ideal generated by the linear network problem and denote this ideal by $\text{Ideal}\left((\mathcal{G}, \mathcal{C})\right)$.

A Little Algebraic Geometry

The algebraic variety associated with Ideal $\left((\mathcal{G},\mathcal{C})\right)$ is given by

$$\begin{aligned}\text{Var}\left((\mathcal{G},\mathcal{C})\right) = \{(a_1,\ldots,a_n) \in \bar{\mathbb{F}}^n \mid \\ f(a_1,\ldots,a_n) = 0 \;\; \forall f \in \text{Ideal}\left((\mathcal{G},\mathcal{C})\right)\}\end{aligned}$$

A Little Algebraic Geometry

The algebraic variety associated with Ideal $((\mathcal{G}, \mathcal{C}))$ is given by

$$\operatorname{Var}((\mathcal{G}, \mathcal{C})) = \{(a_1, \ldots, a_n) \in \bar{\mathbb{F}}^n \mid f(a_1, \ldots, a_n) = 0 \;\; \forall f \in \operatorname{Ideal}((\mathcal{G}, \mathcal{C}))\}$$

Theorem

Let a linear network problem $(\mathcal{G}, \mathcal{C})$ be given. The network problem is solvable if and only if Var$((\mathcal{G}, \mathcal{C}))$ is nonempty and, hence, the ideal Ideal$((\mathcal{G}, \mathcal{C}))$ is a proper ideal of $\bar{\mathbb{F}}[\xi_0, \underline{\xi}]$

Outline

Conclusion

This talk is over.