

On the Rank Metric in Network Error Control Coding

Frank R. Kschischang

*Department of Electrical & Computer Engineering
University of Toronto*


2011 International Symposium on Network Coding
BUPT, Beijing, China

July 25, 2011

Joint work with Danilo Silva and Ralf Kötter


Part I

Error Control



Network Coding


- A new approach to information dissemination over networks
- Essence: packets can be *mixed* with each other (rather than just routed or replicated)
- A higher throughput can be achieved



Linear Network Coding

- Packets are length- m vectors over a finite field \mathbb{F}_q
- Nodes create outgoing packets as \mathbb{F}_q -linear combinations of incoming packets
- Original packets can be recovered by solving a linear system of equations

$$X_i = [X_{i1} \quad \cdots \quad X_{im}]$$



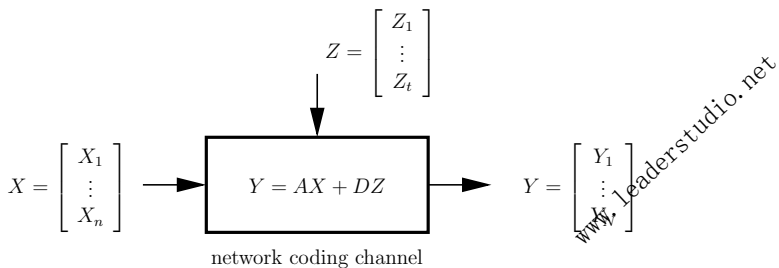

$$\begin{bmatrix} Y_1 \\ \vdots \\ Y_N \end{bmatrix} = \begin{bmatrix} A_{11} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots \\ A_{N1} & \cdots & A_{Nn} \end{bmatrix} \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix}$$

transfer matrix

www.LearnStudy8.net

Linear Network Coding with Errors

- A corrupt packet is modeled as the addition of an error packet to a genuine packet
- Assume that at most t error packets are injected (by an adversary, say)
- The overall network can be viewed as a point-to-point channel



Why Consider Errors?


Possible error sources:

- Random errors that could not be detected at the physical layer
- Corrupt packets injected at the application level by a malicious user

Issue

The very essence of network coding—packet mixing—makes it highly prone to *error propagation*. This essentially rules out classical error correction.

Error Propagation



Two Models

We distinguish between two models for network coding:

- 1 **Coherent**: the network is given, and the local coding coefficients are fixed at design time (so that A is known at all receivers)
- 2 **Noncoherent**: the local coding coefficients are chosen randomly at run time (so that A is not known to the transmitters or receivers)

Under either model, we will assume that an adversary has complete network knowledge.

Deterministic (Coherent) Network Coding:

- 1 N. Cai and R. W. Yeung, "Network coding and error correction," ITW 2002.
- 2 R. W. Yeung \Leftrightarrow N. Cai, "Network error correction, Part I: Basic concepts and upper bounds; Part II: Lower bounds," *Comm. in Inform. and Systems*, 2006.
- 3 R. Matsumoto, "Construction algorithm for network error-correcting codes attaining the Singleton bound," *IEICE Trans. Funda.*, 2007.
- 4 Z. Zhang, "Linear network error correction codes in packet networks," *IEEE Trans. Inf. Theory*, 2008.
- 5 S. Yang, R. W. Yeung, and C. K. Ngai "Refined Coding Bounds and Constructions for Coherent Network Error Correction," *IEEE Trans. Inf. Theory*, 2011.

Random Network Coding:

- 1 S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Médard and M. Effros, "Resilient network coding in the presence of Byzantine adversaries," *IEEE Trans. Inf. Theory*, 2008.
- 2 R. Kötter and F. R. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Trans. Inf. Theory*, 2008.
- 3 D. Silva, F. R. Kschischang, and R. Kötter, "A rank-metric approach to error control in random network coding," *IEEE Trans. Inf. Theory*, 2008.
- 4 A. Montanari and R. Urbanke, "Coding for network coding," 2007, preprint. Available: <http://arxiv.org/abs/0711.3935>

Part II

The Rank Metric

Rank-Metric Codes

Definition

The *rank distance* between matrices $X, X' \in \mathbb{F}_q^{n \times m}$ is defined as $d_R(X, X') = \text{rank}(X - X')$.

Since, for all $X, X', Y \in \mathbb{F}_q^{n \times m}$ we have

- 1 $d_R(X, X') \geq 0$ with equality iff $X = X'$;
- 2 $d_R(X, X') = d_R(X', X)$; and
- 3 $d_R(X, Y) + d_R(Y, X') \geq d_R(X, X')$

the rank distance does indeed qualify as a *bona fide* metric.

A *rank-metric code* $\mathcal{C} \subseteq \mathbb{F}_q^{n \times m}$ is a matrix code used in the context of the rank metric.

The minimum rank distance of \mathcal{C} will be denoted by $d_R(\mathcal{C})$.

Singleton bound for rank metric codes

$$|\mathcal{C}| \leq q^{\max\{n,m\}(\min\{n,m\}-d+1)}$$

for every code $\mathcal{C} \subseteq \mathbb{F}_q^{n \times m}$ with $d_R(\mathcal{C}) = d$.

Codes that achieve this bound are called *maximum-rank-distance* (MRD) codes and linear MRD codes are known to exist for all choices of parameters q , n , m and $d \leq \min\{n, m\}$.

Linearized Polynomials

Let \mathbb{F}_{q^m} be an extension field of \mathbb{F}_q , and recall that \mathbb{F}_{q^m} can be regarded as an m -dimensional vector space over \mathbb{F}_q , i.e., $\mathbb{F}_{q^m} \cong \mathbb{F}_q^m$.

Recall that every function $f : \mathbb{F}_{q^m} \rightarrow \mathbb{F}_{q^m}$ can be represented as a polynomial $f(x) = b_0 + b_1x + \cdots + b_{q^m-1}x^{q^m-1}$. In particular, the \mathbb{F}_q -linear functions can be so represented. (These are the functions, satisfying for all $x, x' \in \mathbb{F}_{q^m}$ and all $a, a' \in \mathbb{F}_q$, $f(ax + ax') = af(x) + a'f(x')$.)

Linearized Polynomials

\mathbb{F}_q -linear functions on \mathbb{F}_{q^m} are represented by so-called *linearized polynomials* w.r.t. \mathbb{F}_q , i.e., polynomials of the form

$$f(x) = f_0x + f_1x^q + f_2x^{q^2} + \cdots + f_{m-1}x^{q^{m-1}} = \sum_{i=0}^{m-1} f_i x^{q^i},$$

where $f_0, \dots, f_{m-1} \in \mathbb{F}_{q^m}$. (Follows from the property that $a^q = a$ if and only if $a \in \mathbb{F}_q$.)

Gabidulin Codes (1985)

Assume $n \leq m$. Let \mathbb{F}_{q^m} be an extension field of \mathbb{F}_q , and let $\theta: \mathbb{F}_{q^m} \rightarrow \mathbb{F}_q^m$ be a vector space isomorphism, where the elements in \mathbb{F}_q^m are regarded as row vectors.

Let $\mathbb{F}_{q,m}[x]$ denote the set of linearized polynomials, i.e., all polynomials of the form $f(x) = \sum_{i=0}^{n-1} f_i x^{q^i}$, where $f_i \in \mathbb{F}_{q^m}$. Let $\alpha_1, \dots, \alpha_n \in \mathbb{F}_{q^m}$ be elements that are linearly independent when regarded as vectors in \mathbb{F}_q^m , and let $0 < d \leq n$.

Evaluation Map

A Gabidulin code $\mathcal{C} \subseteq \mathbb{F}_q^{n \times m}$ is defined as

$$\mathcal{C} = \left\{ [\theta(f(\alpha_1)), \dots, \theta(f(\alpha_n))]^T, f(x) \in \mathbb{F}_{q,m}^{(n-d+1)}[x] \right\}$$

Gabidulin codes are MRD

Theorem

Gabidulin codes are MRD.


Proof: The evaluation map is linear and injective, thus $|\mathcal{C}| = q^{m(n-d+1)}$. We just need to show that the minimum rank of any nonzero codeword is at least d .

The rank of a nonzero codeword $X \in \mathcal{C}$ is equal to the dimension of its image as a linear map; by the rank-nullity theorem we have $\text{rank } X + \dim(\ker X) = n$. However, a polynomial of degree q^{n-d} has at most q^{n-d} zeros; thus $\dim(\ker X) \leq n - d$, from which it follows that $\text{rank } X \geq d$.

Remark: this construction of MRD codes is to linearized polynomials as the original Reed-Solomon construction is to general polynomials.

Part III


Adversarial Channels



Adversarial Channels

An **adversarial channel** is a triple:

$$\left(\underbrace{\mathcal{X}}_{\text{input alphabet}}, \underbrace{\mathcal{Y}}_{\text{output alphabet}}, \underbrace{\{\mathcal{Y}_x \subseteq \mathcal{Y} : x \in \mathcal{X}\}}_{\text{fan-out sets}} \right)$$



When $x \in \mathcal{X}$ is sent, the output y *must* fall in \mathcal{Y}_x (but may be chosen *freely* by an adversary).

Codes for Adversarial Channels

A code \mathcal{C} is a subset of \mathcal{X} . (Note that \mathcal{X} may model multiple uses of a simpler channel.)

A code \mathcal{C} is *unambiguous* if

$$\mathcal{Y}_x \cap \mathcal{Y}_{x'} = \emptyset \text{ for all } x, x' \in \mathcal{C}, x \neq x'$$

Two codewords x and x' are not *distinguishable* if $\mathcal{Y}_x \cap \mathcal{Y}_{x'} \neq \emptyset$.

A *decoder* for \mathcal{C} is any function

$$\hat{x} : \mathcal{Y} \rightarrow \mathcal{C} \cup \{f\},$$

where $f \notin \mathcal{C}$ denotes a decoding failure (detected error).

For example, the *exhaustive* decoder returns

$$\hat{x}(y) = \begin{cases} x & \text{if } y \in \mathcal{Y}_x \text{ and } \nexists x' \in \mathcal{C} \setminus \{x\} \text{ s.t. } y \in \mathcal{Y}_{x'} \\ f & \text{otherwise} \end{cases}$$

Infallible Decoders

If x is sent and $y \in \mathcal{Y}_x$ is received, then \hat{x} is *successful* if

$$\hat{x}(y) = x.$$

A decoder is *infallible* if it never fails, i.e., if it is successful for all $y \in \mathcal{Y}_x$ for all $x \in \mathcal{C}$.

A code \mathcal{C} is unambiguous if and only if it has an infallible decoder.

Ideally we would like an unambiguous \mathcal{C} with

- $|\mathcal{C}|$ as large as possible, and
- an easy-to-implement infallible decoder.

Discrepancy

We would like to consider *families* of adversarial channels, parameterized by a measure of adversarial “effort” t .

We define a *discrepancy* function

$$\Delta : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{N} \triangleq \{0, 1, 2, \dots\}$$

and define, for $t \in \mathbb{N}$,

$$\mathcal{Y}_x(t) = \{y \in \mathcal{Y} : \Delta(x, y) \leq t\}$$

Note that $\mathcal{Y}_x(0) \subseteq \mathcal{Y}_x(1) \subseteq \dots$, i.e., the adversarial channels form a degraded family.

Example: BEC(n)

$$\mathcal{X} = \{0, 1\}^n$$

$$\mathcal{Y} = \{0, 1, \epsilon\}^n$$

$$\Delta(x, y) = \sum_{i=1}^n \Delta(x_i, y_i)$$

where

$$\Delta(x_i, y_i) = \begin{cases} 0 & \text{if } x_i = y_i, \\ 1 & \text{if } y_i = \epsilon, \\ \infty & \text{otherwise} \end{cases}$$

Then $\mathcal{Y}_x(t) = \{y \in \mathcal{Y} : \Delta(x, y) \leq t\}$ consists of all those y that agree with x , except in up to t erased (ϵ) positions.

Minimum-discrepancy Decoder

$$\hat{x}(y) = \operatorname{argmin}_{x \in \mathcal{C}} \Delta(x, y)$$

Min- Δ decoder is infallible provided \mathcal{C} is unambiguous for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$.

(Note that if a decoder is infallible for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$ then it is also infallible for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(s))$ for all $s \leq t$.)

Define the *discrepancy-correction capability* of \mathcal{C} as the largest t for which \mathcal{C} is unambiguous for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$.

Discrepancy-Correction Capability

Let $\tau : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{N}$ be given by

$$\tau(x, x') = \min_{y \in \mathcal{Y}} \max[\Delta(x, y), \Delta(x', y)] - 1,$$

and let

$$\tau(\mathcal{C}) = \min\{\tau(x, x') : x, x' \in \mathcal{C}, x \neq x'\}$$

Theorem

\mathcal{C} is unambiguous for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$ if and only if $t \leq \tau(\mathcal{C})$.

Proof. Take $x, x' \in \mathcal{C}$ with $x \neq x'$.

(\Rightarrow) If \mathcal{C} is unambiguous for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$, then

$\mathcal{Y}_{x_1}(t) \cap \mathcal{Y}_{x_2}(t) = \emptyset$; thus for any $y \in \mathcal{Y}$, either $\Delta(x_1, y) > t$ or $\Delta(x_2, y) > t$ which implies that $\tau(x, x') \geq t$.

(\Leftarrow) If $\tau(\mathcal{C}) \geq t$, then for all $y \in \mathcal{Y}$, then

$\max[\Delta(x, y), \Delta(x', y)] - 1 \geq t$, so $\Delta(x, y) > t$ or $\Delta(x', y) > t$ which implies that $\mathcal{Y}_x(t) \cap \mathcal{Y}_{x'}(t) = \emptyset$. Since x and x' were chosen arbitrarily, \mathcal{C} is unambiguous for $(\mathcal{X}, \mathcal{Y}, \mathcal{Y}_x(t))$.

Remark

$$\tau(x, x') = \min_{y \in \mathcal{Y}} \max[\Delta(x, y), \Delta(x', y)] - 1,$$

Could now define a “distance-like” function given by $2(\tau(x, x') + 1)$ and derive results analogous to classical coding theory; e.g., the “error correction capability is half the minimum distance”.

However, $\tau(x, x')$ is not necessarily very tractable; thus we seek a “simpler” distance measure.

Δ -distance

For $x, x' \in \mathcal{X}$ define

$$\delta(x, x') = \min_{y \in \mathcal{Y}} [\Delta(x, y) + \Delta(x', y)]$$

For example, suppose $\mathcal{X} = \mathcal{Y}$ and Δ is a metric. Then, for any y , $\Delta(x, y) + \Delta(x', y) \geq \Delta(x, x')$ with equality achieved, e.g., if $y = x'$, so in this case Δ -distance corresponds exactly to the metric, i.e., $\delta(x, x') = \Delta(x, x')$.

Discrepancy-correcting capability (2)

Let $\delta(\mathcal{C}) = \min\{\delta(x, x') : x, x' \in \mathcal{C}, x \neq x'\}$

Theorem

$$\tau(\mathcal{C}) \geq \lfloor (\delta(\mathcal{C}) - 1)/2 \rfloor$$


Proof: Take $x, x' \in \mathcal{C}$ with $x \neq x'$, and note that $\max(a, b) \geq \lfloor (a + b + 1)/2 \rfloor$ for all $a, b \in \mathbb{N}$. Now

$$\begin{aligned}\tau(x, x') &= \min_{y \in \mathcal{Y}} \max[\Delta(x, y), \Delta(x', y)] - 1 \\ &\geq \min_{y \in \mathcal{Y}} \lfloor (\Delta(x, y) + \Delta(x', y) + 1)/2 \rfloor - 1 \\ &= \lfloor \frac{\min_{y \in \mathcal{Y}} (\Delta(x, y) + \Delta(x', y) - 1)}{2} \rfloor \\ &= \lfloor (\delta(\mathcal{C}) - 1)/2 \rfloor\end{aligned}$$

Normal Discrepancy Functions

$\Delta : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{N}$ is *normal* if, for all x, x' and all $i \in \{0, \dots, \delta(x, x')\}$ there exists $y \in \mathcal{Y}$ such that

$$\Delta(x, y) = i \text{ and } \Delta(x', y) = \delta(x, x') - i$$



Theorem

If Δ is normal then $\tau(\mathcal{C}) = \lfloor (\delta(\mathcal{C}) - 1)/2 \rfloor$.

BEC(n) revisited

Recall: $\mathcal{X} = \{0, 1\}^n$, $\mathcal{Y} = \{0, 1, \epsilon\}^n$, $\Delta(x, y) = \sum_{i=1}^n \Delta(x_i, y_i)$
where $\Delta(x_i, y_i) \in \{0, 1, \infty\}$.

Here the Δ -distance is

$$\begin{aligned}\delta(x, x') &= \min_{y \in \mathcal{Y}} [\Delta(x, y) + \Delta(x', y)] \\ &= 2d_H(x, x'),\end{aligned}$$

since the min is achieved by the y that erases those positions in which x and x' differ. Thus $\delta(\mathcal{C}) = 2d_H(\mathcal{C})$.

Since the discrepancy function Δ is indeed normal, we have

$$\begin{aligned}\tau(\mathcal{C}) &= \lfloor (\delta(\mathcal{C}) - 1)/2 \rfloor \\ &= \lfloor (2d_H - 1)/2 \rfloor \\ &= d_H - 1\end{aligned}$$

Thus a binary code \mathcal{C} of length n is unambiguous for ρ erasures provided $\rho \leq d_H(\mathcal{C}) - 1$ (as is well known, of course).

Part IV

Coherent Network Coding


Linear Network Coding with Errors

Here $Y = AX + DZ$, where

① $X \in \mathbb{F}_q^{n \times m}$

② $Y \in \mathbb{F}_q^{N \times m}$

③ $A \in \mathbb{F}_q^{N \times n}$ is fixed and known to the receiver and $Z \in \mathbb{F}_q^{t \times m}$ is arbitrarily chosen by an adversary.



We also assume:

④ the adversary is omniscient; in particular the adversary knows A and X

⑤ the matrix $D \in \mathbb{F}_q^{N \times t}$ can be chosen arbitrarily by the adversary.

We may also assume $t < \text{rank}(A)$; otherwise the adversary may always choose $DZ = -AX$, disrupting communication perfectly.

Discrepancy Function

This channel model can be captured by defining:

$$\Delta_A(X, Y) \triangleq \min\{r \in \mathbb{N} : D \in \mathbb{F}_q^{N \times r}, Z \in \mathbb{F}_q^{r \times m}, Y = AX + DZ\}$$

The discrepancy $\Delta_A(X, Y)$ represents the minimum number of error packets that the adversary needs to inject in order to transform input X into output Y , given that the transfer matrix is A .

The minimum-discrepancy decoder becomes

$$\hat{X}(Y) = \operatorname{argmin}_{X \in \mathcal{C}} \Delta_A(X, Y)$$

and the induced Δ -distance is

$$\delta_A(X, X') \triangleq \min_{Y \in \mathbb{F}_q^{N \times m}} [\Delta_A(X, Y) + \Delta_A(X', Y)].$$

Simplifications

Theorem

$$\Delta_A(X, Y) = \text{rank}(Y - AX)$$

Proof: For any triple (r, D, Z) satisfying $Y = AX + DZ$, with $D \in \mathbb{F}_q^{N \times r}$ and $Z \in \mathbb{F}_q^{r \times m}$ we have

$$r \geq \text{rank } Z \geq \text{rank } DZ = \text{rank}(Y - AX),$$

achievable by setting $r = \text{rank}(Y - AX)$ and letting DZ be a full-rank decomposition of $Y - AX$.

Δ -distance

Corollary

$$\delta_A(X, X') = d_R(AX, AX') = \text{rank } A(X - X')$$

Note that $\delta_A(\cdot, \cdot)$ is a metric if and only if A has full column rank — in which case it coincides with the rank metric. (If $\text{rank } A < n$, there exist $X \neq X'$ such that $\Delta_A(X, X') = 0$.)

We can also prove:

Theorem

The discrepancy function $\Delta_A(\cdot, \cdot)$ is normal.

From this we get:

Theorem

A code \mathcal{C} is guaranteed to correct any t packet errors if and only if

$\delta_A(\mathcal{C}) > 2t$, where

$$\delta_A(\mathcal{C}) \stackrel{\Delta}{=} \min\{\text{rank } A(X - X') : X, X' \in \mathcal{C}, X \neq X'\}.$$

Relationship to Rank Distance

Let $\rho = n - \text{rank}(A)$ be the column-rank deficiency of the matrix A . We then have

$$d_R(X, X') - \rho \leq \delta_A(X, X') \leq d_R(X, X').$$

Thus:

Theorem

A code \mathcal{C} with minimum rank distance $d_R(\mathcal{C})$ is guaranteed to correct t packet errors under rank-deficiency ρ if $d_R(\mathcal{C}) > 2t + \rho$.

N.B. This result depends only on ρ and t ; it is independent of the specific transfer matrix A .

Connections to the model of Yeung, et al.

Let $|\mathcal{E}|$ denote the number of edges in the network, and let

$$Y = AX + FE$$

where

- 1 $X \in \mathbb{F}_q^{n \times m}$
- 2 $Y \in \mathbb{F}_q^{N \times m}$
- 3 $A \in \mathbb{F}_q^{N \times n}$ (known and fixed)
- 4 $F \in \mathbb{F}_q^{N \times |\mathcal{E}|}$ (known and fixed), and
- 5 $E \in \mathbb{F}_q^{|\mathcal{E}| \times m}$ is arbitrarily chosen by an adversary provided $\text{wt}(E) \leq t$. The adversary has omniscient knowledge of A , F , and X .

This model may be less “pessimistic” than the previous one, since the adversary is forced to work with a given F .

Discrepancy Function

This channel model can be captured by defining:

$$\Delta_{A,F}(X, Y) \triangleq \min\{\text{wt}(E) : E \in \mathbb{F}_q^{|\mathcal{E}| \times m}, Y = AX + FE\}$$

Can show that:

- 1 this discrepancy measure is normal;
- 2 the decoder of [YY07]¹ is precisely a minimum discrepancy decoder;
- 3 the “network Hamming distance” of [YY07] between two messages is precisely the Δ -distance induced by $\Delta_{A,F}(\cdot, \cdot)$; hence,
- 4 the “unicast minimum distance” of [YY07] is precisely the minimum Δ -distance between codewords;
- 5 this model reduces to the previous one if F can be chosen by the adversary.

¹Shenghao Yang and Raymond W. Yeung, “Characterizations of network error correct/detection and erasure correction,” NetCod, 2007.

Achieving the (Refined) Singleton Bound

From [YeungCai06] (also [YangYeung07])

$$|\mathcal{C}| \leq Q^{n-\rho-\delta_{A,F}(\mathcal{C})+1}$$

where Q is the alphabet size; $Q = q^m$ in our setting. One also obtains

$$|\mathcal{C}| \leq Q^{n-\rho-\delta_A(\mathcal{C})+1}$$

MRD codes with $m \geq n$ achieve

$$\begin{aligned} |\mathcal{C}| &= q^{m(n-d_R(\mathcal{C})+1)} \\ &\geq q^{m(n-\rho-\delta_A(\mathcal{C})+1)} \\ &\geq q^{m(n-\rho-\delta_{A,F}(\mathcal{C})+1)} \end{aligned}$$

Theorem

When $m \geq n$, an MRD code $\mathcal{C} \subseteq \mathbb{F}_q^{n \times m}$ achieves maximum cardinality with respect to both δ_A and $\delta_{A,F}$.

Trading off ρ and t

When $m \geq n$, then $\delta_A(\mathcal{C}) = d_R(\mathcal{C}) - \rho$. Thus:

Theorem


An MRD code $\mathcal{C} \subseteq \mathbb{F}_q^{n \times m}$ with $m \geq n$ is guaranteed to correct t packet errors, under rank deficiency ρ , if and only if $d_R(\mathcal{C}) > 2t + \rho$.

Therefore, we can trade-off a reduced rank-deficiency for improved error-correction capability. Thus we can

- 1 design the network code so that $\text{rank}(A)$ is maximized for each receiver (e.g., by using the LIF algorithm in a network where low min-cut receivers are “augmented” by virtual links, which are later deleted);
- 2 apply an outer MRD code without modification or knowledge of the underlying network code, with the only restriction being $m \geq n$.

Part V

Noncoherent Network Coding



Random Linear Network Coding

- Nodes draw coding coefficients uniformly at random from \mathbb{F}_q
- The transfer matrix will be invertible with high probability if q is sufficiently large
- The transfer matrix can be recorded by appending a header to each original packet

$$\begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix} = \begin{bmatrix} 1 & \cdots & 0 & \boxed{\text{payload 1}} \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & \boxed{\text{payload } n} \end{bmatrix}$$

The Key Idea

In the absence of errors, the transmitter injects X , the receiver collects $Y = AX$. Unfortunately, A is completely unknown to the transmitter and to the receiver (or so we assume).

Q:

What property of X is preserved in AX ?

The Key Idea

In the absence of errors, the transmitter injects X , the receiver collects $Y = AX$. Unfortunately, A is completely unknown to the transmitter and to the receiver (or so we assume).

Q:

What property of X is preserved in AX ?

A:

Left multiplication by A performs *row operations* on $X \Rightarrow$ the rows of AX lie in the **row space** of X .

The Key Idea

In the absence of errors, the transmitter injects X , the receiver collects $Y = AX$. Unfortunately, A is completely unknown to the transmitter and to the receiver (or so we assume).

Q:

What property of X is preserved in AX ?


A:

Left multiplication by A performs *row operations* on $X \Rightarrow$ the rows of AX lie in the **row space** of X .

Thus we may attempt to transmit information via the selection, at the transmitter, of a **vector space** from some appropriate codebook of spaces.


Noncoherent Multi-antenna Channels

The network coding channel resembles the noncoherent multi-antenna channel, **only over \mathbb{F}_q rather than \mathbb{C} .**



Noncoherent Multi-antenna Channels

The network coding channel resembles the noncoherent multi-antenna channel, **only over \mathbb{F}_q rather than \mathbb{C} .**



This subspace approach is inspired by [ZheTse02] (“Communication on the Grassmannian manifold”), where messages are *also* encoded in the choice by the transmitter of an appropriate vector space V .

We will, however, define a different metric on subspaces.

The Operator Channel (A Convenient Abstraction of Random Linear Network Coding)

Let $\mathcal{P}_q(n)$ denote the the set of all subspaces of an n -dimensional vector space over \mathbb{F}_q .

- The transmitter selects a vector space $V \in \mathcal{C}$ from some collection of $\mathcal{C} \subseteq \mathcal{P}_q(n)$ of spaces.
- The transmitter signals this choice by the injection into the network of a basis for V .
- The receiver gathers packets, and forms the vector space U that they span.
- We may write

$$U = \mathcal{H}_k(V) \oplus E,$$

where $\mathcal{H}_k(\cdot)$ is an “**erasure operator**” that projects onto a randomly chosen k -subspace and E denotes an “**error space**” intersecting trivially with V .

Classical Coding Theory

Transmitter: emits a **vector**, e.g., an element of \mathbb{F}_q^n .

Receiver: receives a **vector**, possibly corrupted by noise.

Goal of Code Design: to construct a large collection of **vectors**, well-separated according to some metric (e.g., Hamming distance).

Classical Coding Theory

Transmitter: emits a **vector**, e.g., an element of \mathbb{F}_q^n .

Receiver: receives a **vector**, possibly corrupted by noise.

Goal of Code Design: to construct a large collection of **vectors**, well-separated according to some metric (e.g., Hamming distance).

This Work

Transmitter: emits a **vector space**, e.g., an element of $\mathcal{P}_q(n)$ (the projective space of order n over \mathbb{F}_q).

Receiver: receives a **vector space**, possibly corrupted by noise.

Goal of Code Design: construct a large collection of **vector spaces**, well-separated according to some metric.

Subspace Distance [KK08]

Let A and B be elements of $\mathcal{P}_q(n)$.

Definition

The subspace distance between A and B is defined as

$$d_S(A, B) := \dim(A + B) - \dim(A \cap B).$$

We may also write

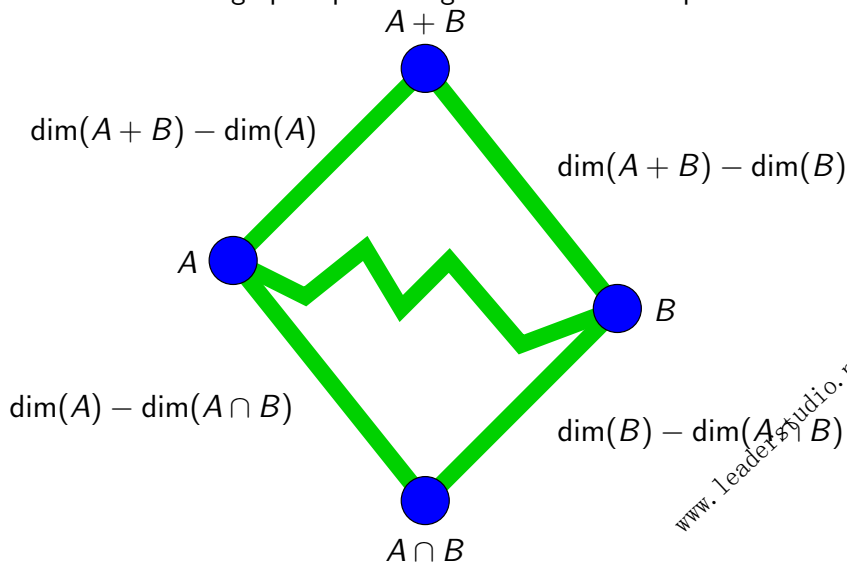
$$\begin{aligned}d_S(A, B) &= \dim(A) + \dim(B) - 2 \dim(A \cap B) \\ &= 2 \dim(A + B) - \dim(A) - \dim(B)\end{aligned}$$

Theorem

The function $d_S(A, B) = \dim(A + B) - \dim(A \cap B)$ is a metric on the space $\mathcal{P}_q(n)$.

Remark:

$d_S(A, B)$ is the length of a geodesic between A and B in the undirected Hasse graph representing the lattice of subspaces.




Injection Distance [SK08]

Definition: (Injection Distance)

$$\begin{aligned}d(A, B) &\triangleq \max\{\dim(\mathbf{A}), \dim(\mathbf{B})\} - \dim(\mathbf{A} \cap \mathbf{B}) \\ &= \dim(\mathbf{A} + \mathbf{B}) - \min\{\dim(\mathbf{A}), \dim(\mathbf{B})\}\end{aligned}$$

$d(\cdot, \cdot)$ is a **metric** that counts the **minimum number of packet injections** required to transform one space to another.



Subspace Distance vs. Injection Distance

- $d_S(A, B)$ is equal to the the minimum number of **insertions and deletions of generators** that are required to transform a basis for A into a basis for B .
(Analogous to Hamming distance in classical coding theory, which is equal to the minimum number of **symbol changes** required to transform a vector A into a vector B .)
- $d(A, B)$ is equal to the minimum number of **packet insertions** needed to transform a basis for A into a basis for B : a *single* packet insertion can simultaneously delete a generator and insert another one.
- If $\dim(A) = \dim(B)$, then $d(A, B) = \frac{1}{2}d_S(A, B)$ (and in general $d(A, B) \geq \frac{1}{2}d_S(A, B)$).
- $d(A, B)$ can be interpreted as a geodesic in a “generalized Grassmann graph.”

Some Definitions

The set of all subspaces of an n -dimensional vector space forms a **projective space** $\mathcal{P}_q(n)$. The set of all ℓ -dimensional subspaces of an n -dimensional vector space is called a **Grassmannian** $\mathcal{G}_q(n, \ell)$.

A **subspace code** $\mathcal{C} \subseteq \mathcal{P}_q(n)$ is a collection of subspaces in $\mathcal{P}_q(n)$. If $\mathcal{C} \subseteq \mathcal{G}_q(n, \ell)$ then \mathcal{C} is a **constant-dimension code** of dimension ℓ .

Definitions (cont'd)

The **minimum distance** of \mathcal{C} is denoted by

$$d(\mathcal{C}) = \min_{X, Y \in \mathcal{C}, X \neq Y} d(X, Y)$$

The **maximum dimension** of \mathcal{C} is denoted by

$$\ell(\mathcal{C}) = \max_{X \in \mathcal{C}} \dim(X)$$

We call \mathcal{C} an $(n, d)_q$ code if $d(\mathcal{C}) = d$; we call \mathcal{C} an $(n, d, \ell)_q$ code if, additionally, $\mathcal{C} \subseteq \mathcal{G}_q(n, \ell)$.

Normalized parameters

rate: $R = \log_q |\mathcal{C}| / (n\ell)$


normalized dimension: $\lambda = \ell(\mathcal{C}) / n$

normalized minimum distance : $\delta = d(\mathcal{C}) / n$

Minimum Distance Decoding

Definition


A *minimum distance decoder* for \mathcal{C} takes the output U of an operator channel and returns a nearest codeword $V \in \mathcal{C}$, i.e., a codeword V satisfying, for all $X \in \mathcal{C}$, $d(U, V) \leq d(U, X)$.



Minimum Distance Decoding

Definition

A *minimum distance decoder* for \mathcal{C} takes the output U of an operator channel and returns a nearest codeword $V \in \mathcal{C}$, i.e., a codeword V satisfying, for all $X \in \mathcal{C}$, $d(U, V) \leq d(U, X)$.



Error-and-Erasure Correcting Capability

Assume we use a code \mathcal{C} for transmission over a random linear network coding channel. Let ρ denote the number of erasures induced by the channel (“rank deficiency,” in the absence of adversarial errors) and let t denote the maximum number of packets that an adversary may inject.

Theorem

A minimum distance decoder for \mathcal{C} will produce the transmitted space V from the received space (for all possible choices of adversarial error) if and only if

$$d(\mathcal{C}) > 2t + \rho.$$

This theorem motivates the construction of codes with large minimum injection distance.

Lifted Rank-Metric Codes

Lifting a Matrix

For a matrix $X \in \mathbb{F}_q^{k \times m}$, let the subspace

$$\Lambda(X) \triangleq \langle [I_{k \times k} \quad X] \rangle \in \mathcal{G}_q(k+m, k)$$

be called the *lifting* of X .

Lifting a Matrix Code

Similarly, for a matrix code $\mathcal{C} \subseteq \mathbb{F}_q^{k \times m}$, let the subspace code

$$\Lambda(\mathcal{C}) \triangleq \{\Lambda(X), X \in \mathcal{C}\}$$

be called the *lifting* of \mathcal{C} .

We have $|\Lambda(\mathcal{C})| = |\mathcal{C}|$ and note that $\Lambda(\mathcal{C})$ is a constant-dimension code.

Relating Rank Distance and Subspace Distance

Theorem

For all $X, X' \in \mathbb{F}_q^{k \times m}$ and all $\mathcal{C} \subseteq \mathbb{F}_q^{k \times m}$,

$$d(\Lambda(X), \Lambda(X')) = d_R(X, X'),$$

$$d(\Lambda(\mathcal{C})) = d_R(\mathcal{C}).$$

Relating Rank Distance and Subspace Distance

Theorem

For all $X, X' \in \mathbb{F}_q^{k \times m}$ and all $\mathcal{C} \subseteq \mathbb{F}_q^{k \times m}$,

$$\begin{aligned}d(\Lambda(X), \Lambda(X')) &= d_R(X, X'), \\d(\Lambda(\mathcal{C})) &= d_R(\mathcal{C}).\end{aligned}$$

Proof. We have

$$\begin{aligned}d(\Lambda(X), \Lambda(X')) &= \dim(\Lambda(X) + \Lambda(X')) - \min\{\dim(\Lambda(X)), \dim(\Lambda(X'))\} \\&= \text{rank} \begin{bmatrix} I & X \\ I & X' \end{bmatrix} - k \\&= \text{rank} \begin{bmatrix} I & X \\ 0 & X' - X \end{bmatrix} - k \\&= \text{rank}(X' - X).\end{aligned}$$

The second statement immediately follows from the first.

Optimality of the Lifting Construction

In particular, let $\mathcal{C} \subseteq \mathbb{F}_q^{k \times (n-k)}$ be an MRD code with $d_{\mathbb{R}}(\mathcal{C}) = d$ and, without loss of generality, let $k \leq n - k$.

Lifted MRD Code

Then $\Lambda(\mathcal{C})$ is an (n, d, k) code with cardinality

$$|\Lambda(\mathcal{C})| = q^{(n-k)(k-d+1)}.$$

Singleton Bound

It is shown in [KK08] that a constant-dimension code \mathcal{C} must satisfy

$$|\mathcal{C}| \leq \begin{bmatrix} n - d + 1 \\ n - k \end{bmatrix}_q < h(q)q^{(n-k)(k-d+1)}$$

where $h(q)$ is a constant depending only on q .

Thus, lifted MRD codes are asymptotically optimal constant-dimension codes.

Conclusions

Error control in linear network coding introduces new and interesting problems in coding theory, in which the rank metric and rank-metric codes play an important role, particularly in the context of adversarial models.

No time to talk about:

- security against wire-tappers in networks with random linear network coding (an Ozarow-Wyner-type coset coding scheme based on lifted rank-metric codes) — see Silva and Kschischang, “Universal Secure Network Coding via Rank-Metric Codes,” *IEEE Trans. Inf. Theory*, Feb. 2011.
- random (non-adversarial) matrix channels — see Silva, Kschischang, Kötter, “Communication over Finite-Field Matrix Channels,” *IEEE Trans. Inf. Theory*, Mar. 2010.