An Algebraic Approach to Physical-Layer Network Coding

Frank R. Kschischang
University of Toronto, Canada

joint work with:
Chen Feng, University of Toronto, Canada
Roberto W. Nóbrega, Federal University of Santa Catarina, Brazil
Danilo Silva, Federal University of Santa Catarina, Brazil

April 18, 2013

WCC, Bergen, Norway

Finite-Field Matrix Channels

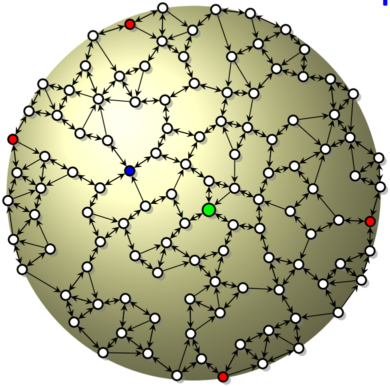

Packet Network

Random Linear Network Coding

- **Transmitter** injects *packets*: vectors from $\mathbb{F}_q^m$, the rows of a matrix X
- Intermediate nodes forward random $\mathbb{F}_q$-linear combinations of packets
- **Errors** may also be injected, which randomly mix with the legitimate packets
- (Each) **receiver** gathers as many packets as possible, forming the rows of matrix Y

At any particular receiver:

$$Y = AX + Z$$

where: X is $n \times m$; Y, Z are $N \times m$; and A is $N \times n$.

A Basic Model

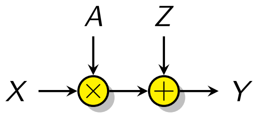

In previous work[1] we considered a basic stochastic linear matrix channel model:

$$Y = AX + Z$$

where

- X and Y are $n \times m$ matrices over $\mathbb{F}_q$;
- A is $n \times n$, nonsingular, drawn uniformly at random;
- Z is $n \times m$ with rank t, drawn uniformly at random;
- X, A, and Z are independent.

[1] D. Silva, K., R. Kötter, "Communication over Finite-Field Matrix Channels," *IEEE Trans. Inf. Theory*, vol. 56, pp. 1296–1305, Mar. 2010.

MAMC: Capacity

Theorem (upper bound)

For $n \leq m/2$,

$$C_{\mathsf{MAMC}} \leq (m-n)(n-t) + \log_q 4(n+1)(t+1).$$

Theorem (lower bound)

Assume $n \leq m$. For any $\epsilon \geq 0$, we have

$$C_{\mathsf{MAMC}} \geq (m-n)(n-t-\epsilon t) - \log_q 4 - \frac{2tnm}{q^{1+\epsilon t}}.$$

These upper and lower bounds match when $q \to \infty$ or $m \to \infty$ (with n/m and t/n fixed).

MAMC: Capacity

Corollary

For large m or large q,

$$C_{\mathrm{MAMC}} \approx (m-n)(n-t).$$

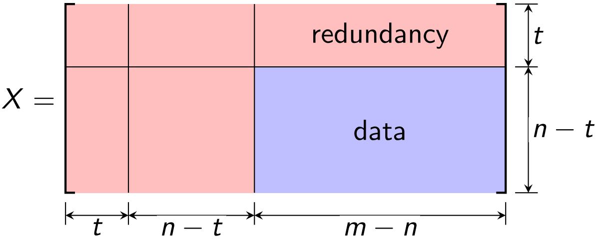

A Simple Coding Scheme

Strategy: Channel Sounding + Error Trapping

Use channel sounding “inside” and error trapping “outside” (but not the opposite!)

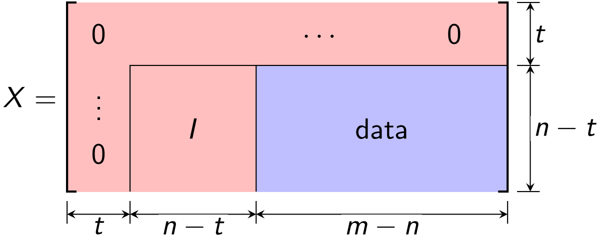

MAMC: A Coding Scheme

First, rewrite the channel model as

$$Y = AX + Z = A(X + A^{-1}Z) = A(X + W), \quad \text{where } W = A^{-1}Z,$$

and suppose a "genie" gives the receiver $X + W$.
Let data matrix D be $(n - t) \times (m - n)$.
We have:

$$X = \begin{bmatrix} 0 & 0 & 0 \\ 0 & I & D \end{bmatrix} \qquad W = \begin{bmatrix} W_1 & W_2 & W_3 \\ W_4 & W_5 & W_6 \end{bmatrix}$$

Assume that $\operatorname{rank} W_1 = t = \operatorname{rank} W$ $(= \operatorname{rank} Z)$. In this case, for some matrix B, we have

$$W = \begin{bmatrix} W_1 & W_2 & W_3 \\ BW_1 & BW_2 & BW_3 \end{bmatrix}$$

Now convert $X + W$ to reduced row echelon (RRE) form:

$$X + W = \begin{bmatrix} W_1 & W_2 & W_3 \\ BW_1 & I + BW_2 & D + BW_3 \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & W_1^{-1}W_2 & W_1^{-1}W_3 \\ BW_1 & I + BW_2 & D + BW_3 \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & W_1^{-1}W_2 & W_1^{-1}W_3 \\ 0 & I & D \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & 0 & \tilde{W}_3 \\ 0 & I & D \end{bmatrix} = \mathrm{RRE}(X + W).$$

But we have Y, not $X + W$!

Now convert $X + W$ to reduced row echelon (RRE) form:

$$X + W = \begin{bmatrix} W_1 & W_2 & W_3 \\ BW_1 & I + BW_2 & D + BW_3 \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & W_1^{-1}W_2 & W_1^{-1}W_3 \\ BW_1 & I + BW_2 & D + BW_3 \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & W_1^{-1}W_2 & W_1^{-1}W_3 \\ 0 & I & D \end{bmatrix}$$

$$\xrightarrow{\text{row op.}} \begin{bmatrix} I & 0 & \tilde{W}_3 \\ 0 & I & D \end{bmatrix} = \mathrm{RRE}(X + W).$$

But we have Y, not $X + W$!

Observation

$Y = A(X + W)$, A is full rank, so Y and $X + W$ have the same row space, which implies that

$$\mathrm{RRE}(Y) = \mathrm{RRE}(X + W).$$

Thus, D is exposed by reducing Y to RRE form!

MAMC: A Coding Scheme

- Decoding amounts to performing full Gaussian elimination on the received matrix Y.

  Complexity: $\mathcal{O}(n^2 m)$ operations in $\mathbb{F}_q$ to recover $(n-t)(m-n)$ symbols. Defining $R = (n-t)(m-t)/mn$, we have a complexity of $\mathcal{O}(n/R)$ operations per decoded symbol.

- The scheme fails if W_1 is not invertible.
  The probability of failure falls exponentially (for fixed m) in the number of bits per field-element, or exponentially (for fixed q) in m (assuming fixed aspect ratio of m/n and fixed t/n).

Theorem

This coding scheme can achieve the capacity of the MAMC when either $q \to \infty$ or $m \to \infty$.

This Talk:

Generalize

from

finite-field matrix channels

to

finite-ring matrix channels.

Why?

This Talk:

Generalize
from
finite-field matrix channels
to
finite-ring matrix channels.

Why?

A: it could be useful for nested-lattice-based **physical-layer** network coding (LNC), a form of compute-and-forward relaying à la

B. Nazer and M. Gastpar, "Compute-and-forward: Harnessing interference through structured codes," *IEEE Trans. Inf. Theory*, vol. 57, pp. 6463–6486, Oct. 2011.

Compute-and-Forward: Nested Lattices

Nested Lattices

Fine lattice Λ, coarse lattice $\Lambda' \subseteq \Lambda$, and lattice quotient Λ/Λ'

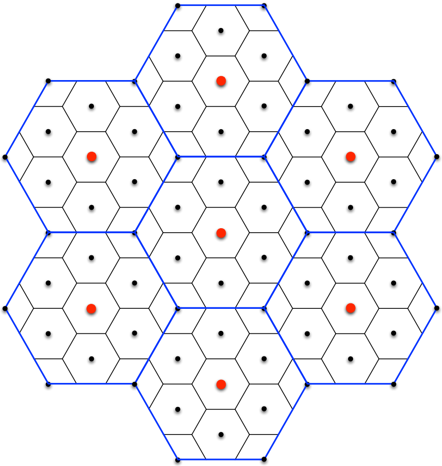

$$\mathbf{G}_\Lambda = \begin{bmatrix} \sqrt{3} & 1 \\ 0 & 2 \end{bmatrix}$$

$$\Lambda = \{\mathbf{r}\mathbf{G}_\Lambda : \mathbf{r} \in \mathbb{Z}^2\}$$

$$\Lambda' = 3\Lambda$$

Compute-and-Forward: Complex Lattices

Complex R-Lattices

Let R be a discrete subring of $\mathbb{C}$ forming a principal ideal domain. Let $N \leq n$. An R-lattice of dimension N in $\mathbb{C}^n$ is defined as the set of all R-linear combinations of N linearly independent vectors, i.e.,

$$\Lambda = \{\mathbf{r}\mathbf{G}_\Lambda : \mathbf{r} \in R^N\},$$

where $\mathbf{G}_\Lambda \in \mathbb{C}^{N\times n}$ is called a generator matrix for Λ.

$R = \mathbb{Z}[\omega] \Rightarrow$ Eisenstein lattices; $R = \mathbb{Z}[i] \Rightarrow$ Gaussian lattices

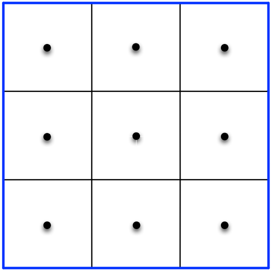

$$\mathbb{Z}[\omega] \triangleq \{a + b\omega : a, b \in \mathbb{Z}, \omega = e^{i2\pi/3}\}$$

$$\mathbb{Z}[i] \triangleq \{a + bi : a, b \in \mathbb{Z}\}$$

$$\Lambda = \mathbb{Z}[i]$$

$$\Lambda' = 3\mathbb{Z}[i]$$

Compute-and-Forward: Structure of Λ/Λ'

Theorem

$$\Lambda/\Lambda' \cong R/\langle \pi_1 \rangle \times \cdots \times R/\langle \pi_k \rangle$$

for some nonzero, non-unit $\pi_1, \ldots, \pi_k \in R$ such that $\pi_1 \mid \cdots \mid \pi_k$. Moreover, there exists a surjective R-module homomorphism $\varphi : \Lambda \to R/\langle \pi_1 \rangle \times \cdots \times R/\langle \pi_k \rangle$ whose kernel is Λ'.

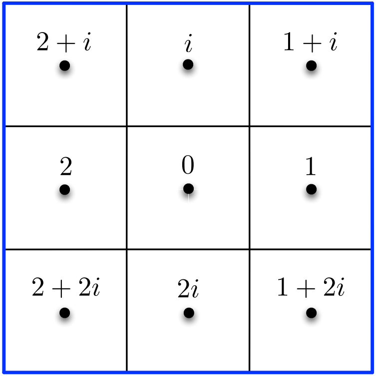

$$\Lambda/\Lambda' \cong \mathbb{Z}[i]/\langle 3 \rangle$$

$$\varphi(a+bi) = (a+bi) \bmod 3$$

$$\varphi^{-1}(c+di) = (c+di) + \Lambda'$$

Compute-and-Forward: Architecture

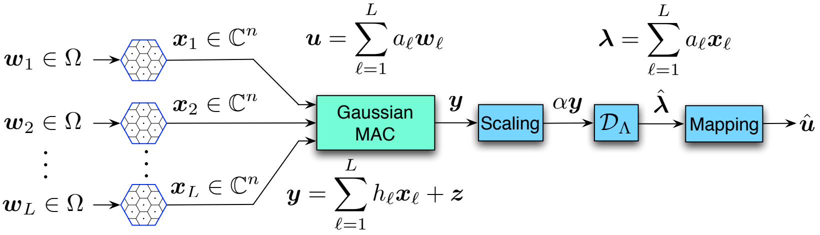

$R/\langle \pi_1 \rangle \times \cdots \times R/\langle \pi_k \rangle$ is the message space Ω

Encoding

Transmitter ℓ sends $\mathbf{x}_\ell \in \Lambda$, a coset representative of $\varphi^{-1}(\mathbf{w}_\ell)$

Decoding

Receiver first recovers $\sum_\ell a_\ell \mathbf{x}_\ell$ from $\alpha \mathbf{y}$;
Receiver then maps $\sum_\ell a_\ell \mathbf{x}_\ell$ onto $\sum_\ell a_\ell \mathbf{w}_\ell$ via φ

Remark: $\alpha \mathbf{y} - \sum_\ell a_\ell \mathbf{x}_\ell = \sum_\ell (\alpha h_\ell - a_\ell) \mathbf{x}_\ell + \alpha \mathbf{z}$ "effective noise"

Construction Examples

Example 1: [Ordentlich, Zhan, Erez, Gastpar, Nazer, ISIT'11]

- Λ is obtained using Construction A applied to binary $(n = 64800, k = 54000)$ LDPC code C, with mod-4 shaping:

$$\Lambda = C + 2\mathbb{Z}^n, \quad \Lambda' = 4\mathbb{Z}^n.$$

- Induced message space: $\mathbb{Z}_4^{54000} \times \mathbb{Z}_2^{10800}$

Example 2: Turbo Lattices [Sakzad, Sadeghi, Panario, Allerton'10]

- Λ is obtained using Construction D applied to nested turbo codes $C_2 : (n = 10131, k_2 = 3377)$ and $C_1 : (n = 10131, k_1 = 5065)$;

$$\Lambda = C_2 + 2C_1 + 4\mathbb{Z}^n, \quad \Lambda' = 4\mathbb{Z}^n.$$

- Induced message space: $\mathbb{Z}_4^{3377} \times \mathbb{Z}_2^{1688}$

In general, for most practical constructions, we have

$$\Omega = R/\langle \pi^{t_0} \rangle \times \cdots \times R/\langle \pi^{t_{m-1}} \rangle, \ t_0 \geq \ldots \geq t_{m-1}.$$

Much Ongoing Work:

B. Nazer and M. Gastpar, “Compute-and-forward: Harnessing interference through structured codes,” *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6463–6486, Oct. 2011.

M. P. Wilson, K. Narayanan, H. D. Pfister, and A. Sprintson, “Joint physical layer coding and network coding for bidirectional relaying,” *IEEE Trans. Inf. Theory*, vol. 56, no. 11, pp. 5641–5654, Nov. 2010.

N. E. Tunali, K. R. Narayanan, J. J. Boutros, and Y.-C. Huang, “Lattices over Eisenstein integers for compute-and-forward,” in *Proc. 2012 Allerton Conf. Commun., Control, and Comput.*, Monticello, IL, Oct. 2012, pp. 33–40.

S. Qifu and J. Yuan, “Lattice network codes based on Eisenstein integers,” in *Proc. 2012 IEEE Int. Conf. on Wireless and Mobile Comput.*, Barcelona, Spain, Oct. 2012, pp. 225–231.

A. Osmane and J.-C. Belfiore, “The compute-and-forward protocol: implementation and practical aspects,” 2011.

S. Gupta and M. A. Vázquez-Castro, “Physical-layer network coding based on integer-forcing precoded compute-and-forward,” 2013.

Chain Rings,
Modules,
Matrices

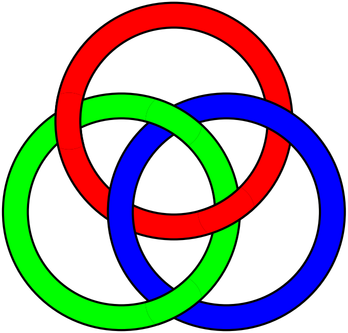

Commutative Rings with Identity $1 \neq 0$

- Ideals in a ring can be partially ordered by subset inclusion.
- The resulting poset is called the lattice of ideals of the ring.

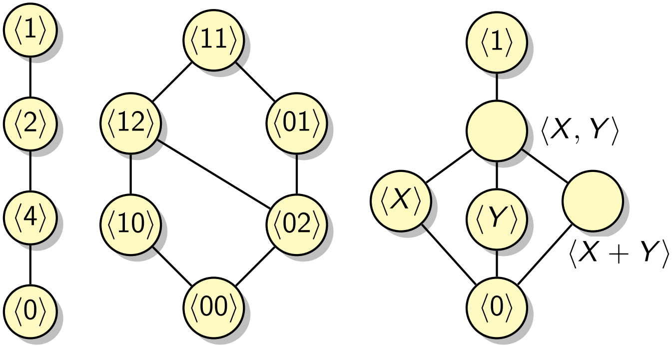

$\mathbb{Z}_8$ $\qquad$ $\mathbb{Z}_2 \times \mathbb{Z}_4$ $\qquad$ $\mathbb{Z}_2[X, Y]/\langle X, Y\rangle^2$

Chain ring: ideals are linearly ordered. Ex: $\mathbb{Z}_8$.

Principal ideal ring: every ideal gen. by 1 element. Ex: $\mathbb{Z}_8$, $\mathbb{Z}_2 \times \mathbb{Z}_4$.

Local ring: unique maximal proper ideal. Ex: $\mathbb{Z}_8$, $\mathbb{Z}_2[X, Y]/\langle X, Y\rangle^2$.

Finite Rings: Important Facts

Proposition

If R is a ring and N is a **maximal** ideal of R, then R/N is a **field**.

This is called a residue field.

Proposition

A finite ring is a **chain** ring if and only if it is both **local** and **principal**.

Proposition

Every finite **principal ideal ring** is a product of finite **chain rings**.

Finite Chain Rings: The Ideals

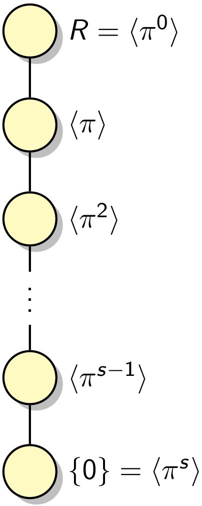

Let R be a finite chain ring, where

- $\langle \pi \rangle$ is the unique maximal ideal,
- q is the order of the residue field,
- s is the number of proper ideals.

Proposition

The lattice of ideals of R is

$$R \supset \langle \pi \rangle \supset \langle \pi^2 \rangle \supset \cdots \supset \langle \pi^{s-1} \rangle \supset \langle \pi^s \rangle = \{0\}.$$

We have $|\langle \pi^i \rangle| = q^{s-i}$; in particular $|R| = q^s$.

Notation: (q, s) chain ring.

Finite Chain Rings: Examples

The following are two non-isomorphic ($q = 2, s = 2$) chain rings.

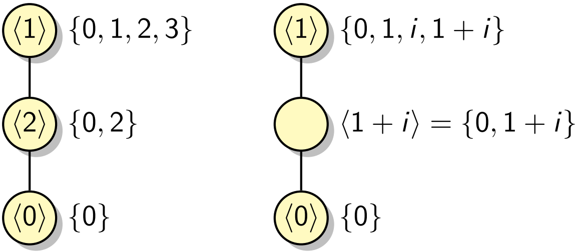

$\mathbb{Z}_4 = \{0, 1, 2, 3\}$ $\qquad$ $\mathbb{Z}_2[i] = \{0, 1, i, 1 + i\}$

In other words, specifying q and s does not uniquely specify the chain ring.

Finite Chain Rings: The π-adic Decomposition

Let R be a (q, s) chain ring.

Proposition

Fix the following:

- $\pi \in R$, a generator for the maximal ideal $\langle \pi \rangle$.
- $\mathcal{R}(R, \pi)$, a complete set of residues with respect to π.

Then every element $r \in R$ can be written uniquely as

$$r = r_0 + r_1\pi + r_2\pi^2 + \cdots + r_{s-1}\pi^{s-1}$$

where $r_i \in \mathcal{R}(R, \pi)$.

This is known as the π-adic decomposition.

Element Degree

Definition

The degree, $\deg(r)$, of a nonzero element $r \in R^*$, where

$$r = r_0 + r_1\pi + \cdots + r_{s-1}\pi^{s-1},$$

is defined as the *least* index j for which $r_j \neq 0$.

- by convention, $\deg(0) = s$
- units have degree zero
- elements of the same degree are associates
- a divides b if and only if $\deg(a) \leq \deg(b)$
- $\deg(a+b) \geq \min\{\deg(a), \deg(b)\}$

Shapes

An *s-shape* $\mu = (\mu_1, \mu_2, \ldots, \mu_s)$ is a sequence of non-decreasing non-negative integers, i.e., $0 \leq \mu_1 \leq \mu_2 \leq \cdots \leq \mu_s$. We denote by $|\mu|$ the sum of its components, i.e., $|\mu| = \sum_{i=1}^{s} \mu_i$.

Example: $\mu = (4, 6, 8)$

```
* * * *
* * * * * *
* * * * * * * *
```

$|(4, 6, 8)| = 18$

For convenience, we will sometimes identify the integer t with the s-shape $(t, \ldots, t)$.

An s-shape $\kappa = (\kappa_1, \ldots, \kappa_s)$ is said to be a subshape of $\mu = (\mu_1, \ldots, \mu_s)$, written $\kappa \preceq \mu$, if $\kappa_i \leq \mu_i$ for all $i = 1, \ldots, s$.

```
* * * *
* * * * * *
* * * * * * * *
```

$(4, 4, 5) \preceq (4, 6, 8)$

From Shape to Module

When R is a finite chain ring, an R-module is always isomorphic to a direct product of various ideals of R; this structure can be described by a *shape*.

Definition

Let R be a (q,s) chain ring with maximal ideal $\langle\pi\rangle$. For any s-shape μ, we define the R-module R^μ as

$$R^\mu \triangleq \underbrace{\langle 1\rangle \times \cdots \times \langle 1\rangle}_{\mu_1} \times \underbrace{\langle \pi\rangle \times \cdots \times \langle \pi\rangle}_{\mu_2-\mu_1} \times \cdots \times \underbrace{\langle \pi^{s-1}\rangle \times \cdots \times \langle \pi^{s-1}\rangle}_{\mu_s-\mu_{s-1}}.$$

$R^{(\mu_1,\ldots,\mu_s)}$ is a collection of μ_s-tuples over R, whose π-adic coordinate array must satisfy degree constraints specified by $(\mu_1,\ldots,\mu_s)$.

Note that $|R^\mu| = q^{|\mu|}$.

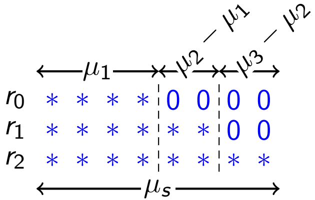

$s = 3,\ \mu = (4,6,8)$

From Module to Shape

Conversely, we have the following theorem (see, e.g., [HL00][2]).

Theorem

For any finite R-module M over a (q, s) chain ring R, there is a unique s-shape μ such that $M \cong R^{\mu}$.

- We call the unique shape μ associated with a module M the shape of M, and write $\mu = \text{shape}\, M$.
- If M' is a submodule of M, then $\text{shape}\, M' \preceq \text{shape}\, M$, i.e., the shape of a submodule is a subshape of the module.

For example, the module spanned by 1111 and 0022 over $\mathbb{Z}_8$ has shape (1,2,2). This module contains 2^5 4-tuples, and is isomorphic to $\langle 1 \rangle \times \langle 2 \rangle$.

[2] T. Honold and I. Landjev, "Linear Codes over Finite Chain Rings," *The Electronic J. of Combinatorics*, vol. 7, 2000.

Counting Submodules

It is also known [HL00] that the number of submodules of R^μ whose shape is κ is given by

$$\left\llbracket \begin{matrix} \mu \\ \kappa \end{matrix} \right\rrbracket_q = \prod_{i=1}^{s} q^{(\mu_i - \kappa_i)\kappa_{i-1}} \begin{bmatrix} \mu_i - \kappa_{i-1} \\ \kappa_i - \kappa_{i-1} \end{bmatrix}_q, \tag{1}$$

where

$$\begin{bmatrix} m \\ k \end{bmatrix}_q \triangleq \prod_{i=0}^{k-1} \frac{q^m - q^i}{q^k - q^i}$$

is the Gaussian coefficient.

In particular, when the chain length $s = 1$, R becomes the finite field $\mathbb{F}_q$ of q elements, and $\left\llbracket \begin{smallmatrix} \mu \\ \kappa \end{smallmatrix} \right\rrbracket_q$ becomes $\left[\begin{smallmatrix} \mu_1 \\ \kappa_1 \end{smallmatrix} \right]_q$, which is the number of κ_1-dimensional subspaces of $\mathbb{F}_q^{\mu_1}$.

Matrices over Finite Chain Rings

Notation for matrices:

- $R^{n\times m}$: the set of all $n \times m$ matrices with entries from ring R.
- $U \in R^{n\times n}$ is invertible if $UV = VU = I_n$ for some $V \in R^{n\times n}$, where I_n denotes the $n \times n$ identity matrix. The set of invertible matrices in $R^{n\times n}$ forms the general linear group $\mathrm{GL}_n(R)$ under multiplication.
- $A, B \in R^{n\times m}$ are left-equivalent if there exists a matrix $U \in \mathrm{GL}_n(R)$ such that $UA = B$.
- $A, B \in R^{n\times m}$ are equivalent if there exist matrices $U \in \mathrm{GL}_n(R)$ and $V \in \mathrm{GL}_m(R)$ such that $UAV = B$.
- $D \in R^{n\times m}$ is a diagonal matrix if $D[i,j] = 0$ whenever $i \neq j$. A diagonal matrix need not be square.

$$\begin{bmatrix} * & 0 & 0 & 0 \\ 0 & * & 0 & 0 \\ 0 & 0 & * & 0 \end{bmatrix} \qquad \begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \\ 0 & 0 & 0 \end{bmatrix}$$

Smith Normal Form

Definition

A diagonal matrix $D = \operatorname{diag}(d_1, \ldots, d_r) \in R^{n \times m}$ is called a Smith normal form of $A \in R^{n \times m}$, if D is equivalent to A and $d_1 \mid d_2 \mid \cdots \mid d_r$ in R, where $r = \min\{n, m\}$.

Every matrix over a PIR (in particular, a finite chain ring) has a Smith normal form whose diagonal entries are unique up to equivalence of associates.

Over $R = \mathbb{Z}_8$

$$A = \begin{bmatrix} 4 & 6 & 2 & 1 \\ 0 & 0 & 0 & 2 \\ 2 & 4 & 6 & 1 \\ 2 & 0 & 2 & 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 2 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}}_{U} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}}_{S} \underbrace{\begin{bmatrix} 0 & 2 & 2 & 1 \\ 1 & 1 & 2 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{V}$$

with invertible U and V. Since $1 \mid 2 \mid 4 \mid 0$ in $\mathbb{Z}_8$, S is the Smith normal form of A.

Row and Column Span

- For $A \in R^{n \times m}$, denote by $\operatorname{row} A$ and $\operatorname{col} A$ the row span and column span of A, respectively.
- From the Smith normal form, it is easy to see that $\operatorname{row} A \cong \operatorname{col} A$.
- Two matrices $A, B \in R^{n \times m}$ are left-equivalent if and only if $\operatorname{row} A = \operatorname{row} B$.
- Two matrices $A, B \in R^{n \times m}$ are equivalent if and only if $\operatorname{row} A \cong \operatorname{row} B$.

Shape of a Matrix

Definition

The *shape* of a matrix A is defined as the shape of the row span of A, i.e.,

$$\operatorname{shape} A = \operatorname{shape}(\operatorname{row} A).$$

Clearly, $\operatorname{shape} A = \operatorname{shape}(\operatorname{col} A)$.
$\operatorname{shape} A = \mu$ if and only if the Smith normal form of A is given by

$$\operatorname{diag}(\underbrace{1, \dots, 1}_{\mu_1}, \underbrace{\pi, \dots, \pi}_{\mu_2 - \mu_1}, \dots, \underbrace{\pi^{s-1}, \dots, \pi^{s-1}}_{\mu_s - \mu_{s-1}}, \underbrace{0, \dots, 0}_{r - \mu_s}),$$

where $r = \min\{n, m\}$.
A matrix $U \in R^{n \times n}$ is invertible if and only if $\operatorname{shape} U = (n, \dots, n)$.

Example

If A has Smith normal form $D = \operatorname{diag}(1, 2, 4, 0)$ over $\mathbb{Z}_8$ then $\operatorname{shape} A = (1, 2, 3)$.

Properties of Matrix Shape

Let $A \in R^{n \times m}$ and $B \in R^{m \times k}$. Then

- shape $A =$ shape A^T, where A^T is the transpose of A.
- For any $P \in \mathrm{GL}_n(R), Q \in \mathrm{GL}_m(R)$, shape $A =$ shape PAQ.
- shape $AB \preceq$ shape A, shape $AB \preceq$ shape B.
- For any submatrix C of A, shape $C \preceq$ shape A.

Row Canonical Form

Let R be a (q, s) chain ring with maximal ideal $\langle \pi \rangle$, fixing a complete set of residues $\mathcal{R}(R, \pi)$ (including 0), and for $1 < \ell < s$, fixing

$$\mathcal{R}(R, \pi^\ell) = \left\{ \sum_{i=0}^{\ell-1} a_i \pi^i : a_0, \ldots, a_{\ell-1} \in \mathcal{R}(R, \pi) \right\}.$$

Example: $R = \mathbb{Z}_8$

If $R = \mathbb{Z}_8$, with $\pi = 2$, we might fix $\mathcal{R}(R, 2) = \{0, 1\}$, so that $\mathcal{R}(R, 4) = \{0, 1, 2, 3\}$.

Row Canonical Form (cont'd)

In a matrix A:

- The element $A[i, j]$ occurs above $A[i', j']$ if $i < i'$. (Equivalently, $A[i', j']$ occurs below $A[i, j]$.)
- The element $A[i, j]$ occurs earlier than $A[i', j']$ if $j < j'$. (Equivalently, $A[i', j']$ occurs later than $A[i, j]$.)
- The first element in row i with property P occurs earlier than any other element in row i with property P.
- The pivot of a nonzero row of A is the first entry among the entries having least degree in that row. For example, the pivot of [0 4 6 2] over $\mathbb{Z}_8$ is the element 6.

Row Canonical Form (cont'd)

Definition

A matrix A is in *row canonical form* if it satisfies the following conditions.

1. Nonzero rows of A are above any zero rows.
2. If A has two pivots of the same degree, the one that occurs earlier is above the one that occurs later. If A has two pivots of different degree, the one with smaller degree is above the one with larger degree.
3. Every pivot is of the form π^ℓ for some $\ell \in \{0, \ldots, s-1\}$.
4. For every pivot (say π^ℓ), all entries below and in the same column as the pivot are zero, and all entries above and in the same column as the pivot are elements of $\mathcal{R}(R, \pi^\ell)$.

For example, over $\mathbb{Z}_8$,

$$A = \begin{bmatrix} 0 & 2 & 0 & \bar{1} \\ \bar{2} & 2 & 0 & 0 \\ 0 & 0 & \bar{2} & 0 \\ 0 & \bar{4} & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

is in row canonical form.

Basic Facts

Let $A \in R^{n \times m}$ be a matrix in row canonical form, let p_k be the pivot of the kth row, let c_k be the index of the column containing p_k. (If the kth row is zero, let $p_k = 0$ and $c_k = 0$.) Let $d_k = \deg(p_k)$, and let $w = (w_1, \ldots, w_m)$ be an arbitrary element of row A.

1. Any column of A contains at most one pivot.
2. If A has more than one row, deleting a row of A results in a matrix also in row canonical form.
3. $i \geq k$ implies $\deg(A[i, j]) \geq d_k$.
4. ($i \geq k$ and $j < c_k$) or ($i > k$ and $j \leq c_k$) implies $\deg(A[i, j]) > d_k$.
5. p_1 divides $w_1, w_2, \ldots, w_m$.
6. $j < c_1$ implies $\deg(w_j) > d_1$.

Reduction to Row Canonical Form

PivotSelection: given a submatrix, return the row and column index of the earliest occurring pivot of least possible degree; otherwise declare the submatrix to be zero.

Given a matrix A:

- Step $k = 1$: apply PivotSelection to A; move the selected row to row 1, normalize (make sure the first pivot is of the form π^ℓ), and cancel all elements below the pivot (which must all be multiples of the first pivot). Call the resulting matrix A_1, and increment k.
- For $k \geq 2$, apply PivotSelection to the rows of A_{k-1}, excluding the first $k-1$ rows. If no pivot can be found, stop; otherwise, move the selected row to row k, normalize to π^ℓ, cancel all elements below the pivot, and reduce all elements above the pivot to elements of $\mathcal{R}(R, \pi^\ell)$. Call the resulting matrix A_k, and increment k.

Row Canonical Form (cont'd)

Theorem

For any $A \in R^{n\times m}$, the algorithm described above computes a row canonical form of A.

Theorem

For any $A \in R^{n\times m}$, the row canonical form of A is unique.

Example:

$$A = \begin{bmatrix} 4 & 6 & 2 & \bar{1} \\ 0 & 0 & 0 & 2 \\ 2 & 4 & 6 & 1 \\ 2 & 0 & 2 & 1 \end{bmatrix} \to A_1 = \begin{bmatrix} 4 & 6 & 2 & 1 \\ 0 & 4 & 4 & 0 \\ \bar{6} & 6 & 4 & 0 \\ 6 & 2 & 0 & 0 \end{bmatrix} \to$$

$$A_1' = \begin{bmatrix} 4 & 6 & 2 & 1 \\ \bar{2} & 2 & 4 & 0 \\ 0 & 4 & 4 & 0 \\ 6 & 2 & 0 & 0 \end{bmatrix} \to A_2 = \begin{bmatrix} 0 & 2 & 2 & 1 \\ 2 & 2 & 4 & 0 \\ 0 & \bar{4} & 4 & 0 \\ 0 & 4 & 4 & 0 \end{bmatrix} \to$$

$$A_3 = \begin{bmatrix} 0 & 2 & 2 & \bar{1} \\ \bar{2} & 2 & 4 & 0 \\ 0 & \bar{4} & 4 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$ which is in row canonical form.

Matrix Shape via Row Canonical Form

Let B be the row canonical form of $A \in R^{n \times m}$ with k nonzero rows. Let p_i be the pivot in the ith row of B, where $i \in \{1, \ldots, k\}$. Let $r = \min\{n, m\}$. Clearly, $k \leq r$. Then the Smith normal form of A is given by

$$\operatorname{diag}(p_1, \ldots, p_k, \underbrace{0, \ldots, 0}_{r-k}) \in R^{n \times m},$$

from which the shape of A is readily available.
Example:

$$A = \begin{bmatrix} 4 & 6 & 2 & 1 \\ 0 & 0 & 0 & 2 \\ 2 & 4 & 6 & 1 \\ 2 & 0 & 2 & 1 \end{bmatrix} \to B = \begin{bmatrix} 0 & 2 & 2 & 1 \\ 2 & 2 & 4 & 0 \\ 0 & 4 & 4 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

over $\mathbb{Z}_8$. Since B is the row canonical form of A, we see that the Smith normal form is $\operatorname{diag}(1, 2, 4, 0)$, and hence shape $A = (1, 2, 3)$.

π-adic Decomposition

Let $R^{n\times\mu}$ denote the set of matrices in $R^{n\times m}$ whose rows are elements of R^{μ}. Every matrix X in $R^{n\times\mu}$ decomposes according to its π-adic decomposition as

$$X = X_0 + \pi X_1 + \cdots + \pi^{s-1} X_{s-1},$$

with each auxiliary matrix X_i $(i = 0, \ldots, s-1)$ satisfying:

1. $X_i[1{:}n, 1{:}\mu_{i+1}]$ is an arbitrary matrix over $\mathcal{R}(R, \pi)$, and
2. all other entries in X_i are zero.

Example: $n = 6$, $\mu = (4, 6, 8)$.

$* \ * \ * \ *$

$* \ * \ * \ * \ * \ *$

$* \ * \ * \ * \ * \ * \ * \ *$

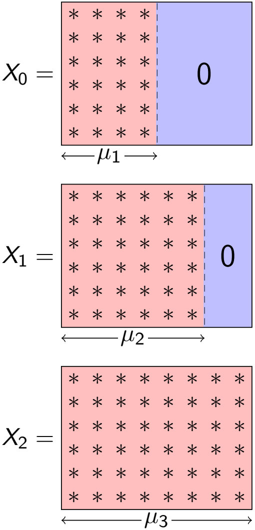

Row Canonical Forms in $\mathcal{T}_\kappa(R^{n\times\mu})$

Let $\mathcal{T}_\kappa(R^{n\times\mu})$ denote the set of matrices in $R^{n\times\mu}$ whose shape is κ, where $\kappa \preceq n$ and $\kappa \preceq \mu$.
The row canonical forms in $\mathcal{T}_\kappa(R^{n\times\mu})$ are in one-to-one correspondence with the submodules of R^μ having shape κ; thus there are $\left[\!\left[\begin{smallmatrix}\mu\\ \kappa\end{smallmatrix} \right]\!\right]_q$ such row canonical forms.

Example:

Let $R = \mathbb{Z}_4$, and let $n = 2$, $\mu = (2,3)$, $\kappa = (1,2)$. Then $\left[\!\left[\begin{smallmatrix}\mu\\ \kappa\end{smallmatrix} \right]\!\right]_q = 18$.
These 18 row canonical forms can be classified into 4 categories based on the positions of their pivots:

$$\underbrace{\begin{bmatrix}1 & * & * \\ 0 & 2 & *\end{bmatrix}}_{8} \quad \underbrace{\begin{bmatrix}0 & 1 & * \\ 2 & 0 & *\end{bmatrix}}_{4} \quad \underbrace{\begin{bmatrix}1 & * & 0 \\ 0 & 0 & 2\end{bmatrix}}_{4} \quad \underbrace{\begin{bmatrix}* & 1 & 0 \\ 0 & 0 & 2\end{bmatrix}}_{2}$$

Clearly, the first category, whose pivots occur as early as possible, contains a significant portion of all possible row canonical forms.

Principal RCFs — The "Thick Cell"

Definition

A row canonical form in $\mathcal{T}_\kappa(R^{n\times\mu})$ is called *principal* if its diagonal entries $d_1, d_2, \ldots, d_r$ $(r = \min\{n, m\})$ have the following form:

$$d_1, \ldots, d_r = \underbrace{1, \ldots, 1}_{\kappa_1}, \underbrace{\pi, \ldots, \pi}_{\kappa_2-\kappa_1}, \ldots, \underbrace{\pi^{s-1}, \ldots, \pi^{s-1}}_{\kappa_s-\kappa_{s-1}}, \underbrace{0, \ldots, 0}_{r-\kappa_s}.$$

All principal RCFs in $\mathcal{T}_\kappa(R^{n\times\mu})$ can be constructed via a π-adic decomposition:

$$X_0 = \begin{array}{c} \xleftarrow{\ \kappa_1\ }\!\!\xrightarrow{} \\ \left[\begin{array}{cc:cc:cccc} 1 & & * & * & & & & \\ & 1 & * & * & & & & \\ \hdashline & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \end{array}\right] \\ \xleftarrow{\quad \mu_1 \quad}\!\!\xrightarrow{} \end{array} \quad X_1 = \begin{array}{c} \xleftarrow{\ \kappa_2\ }\!\!\xrightarrow{} \\ \left[\begin{array}{ccc:ccc:cc} 0 & & & * & * & * & & \\ & 0 & & * & * & * & & \\ & & 1 & * & * & * & & \\ \hdashline & & & & & & & \\ & & & & & & & \\ & & & & & & & \end{array}\right] \\ \xleftarrow{\quad \mu_2 \quad}\!\!\xrightarrow{} \end{array} \quad X_2 = \begin{array}{c} \xleftarrow{\ \kappa_3\ }\!\!\xrightarrow{} \\ \left[\begin{array}{cccc:cccc} 0 & & & & * & * & * & * \\ & 0 & & & * & * & * & * \\ & & 0 & & * & * & * & * \\ & & & 1 & * & * & * & * \\ \hdashline & & & & & & & \\ & & & & & & & \end{array}\right] \\ \xleftarrow{\quad \mu_3 \quad}\!\!\xrightarrow{} \end{array}$$

Illustration of the construction of principal row canonical forms for $\mathcal{T}_\kappa(R^{n\times\mu})$ with $s = 3$, $n = 6$, $\mu = (4, 6, 8)$, and $\kappa = (2, 3, 4)$.

Counting Principal RCFs in $\mathcal{T}_\kappa(R^{n\times\mu})$

Note that the number of principal row canonical forms in $\mathcal{T}_\kappa(R^{n\times\mu})$ is

$$P_q(\mu,\kappa) = q^{\sum_{i=1}^{s} \kappa_i(\mu_i-\kappa_i)}.$$

The number of row canonical forms in $\mathcal{T}_\kappa(R^{n\times\mu})$ in total is

$$\left[\!\left[\begin{matrix} \mu \\ \kappa \end{matrix} \right]\!\right]_q = \prod_{i=1}^{s} q^{(\mu_i-\kappa_i)\kappa_{i-1}} \begin{bmatrix} \mu_i - \kappa_{i-1} \\ \kappa_i - \kappa_{i-1} \end{bmatrix}_q$$

Since $q^{k(m-k)} \leq \begin{bmatrix} m \\ k \end{bmatrix}_q < 4q^{k(m-k)}$, we have

$$1 \leq \frac{\left[\!\left[\begin{matrix} \mu \\ \kappa \end{matrix} \right]\!\right]_q}{P_q(\mu,\kappa)} < 4^s,$$

i.e., the number of principal RCFs in $\mathcal{T}(R^{n\times\mu})$ grows at the same rate as the number of RCFs in total.

Counting All Matrices in $\mathcal{T}_\kappa(R^{n\times\mu})$

We can partition the matrices in $\mathcal{T}_\kappa(R^{n\times\mu})$ based on their row canonical forms: two matrices belong to the same class if and only if they have the same row canonical form.

- The number of classes is $\left[\!\left[{\mu \atop \kappa}\right]\!\right]_q$.
- The number of matrices in each class is

$$|R^{n\times\kappa}| \prod_{i=0}^{\kappa_s-1}(1-q^{i-n}) = q^{n|\kappa|} \prod_{i=0}^{\kappa_s-1}(1-q^{i-n}) =$$

- It follows that

$$|\mathcal{T}_\kappa(R^{n\times\mu})| = q^{n|\kappa|} \prod_{i=0}^{\kappa_s-1}(1-q^{i-n}) \left[\!\left[{\mu \atop \kappa}\right]\!\right]_q .$$

Matrix Channels over Finite Chain Rings

Let R be a (q,s) chain ring, and let μ be an s-shape. We think of R^μ as the "packet space" associated with a network. The length, m, of each packet is given by μ_s.

- The transmitter sends n packets, each constrained to be an element of R^μ. These form the rows of the transmitted matrix $X \in R^{n\times\mu}$.
- The receiver gathers N packets, each also an element of R^μ. These form the rows of the received matrix $Y \in R^{N\times\mu}$.
- Noise is modelled by the injection of t packets into the network, each also an element of R^μ. These form the rows of the noise matrix $Z \in R^{t\times\mu}$.

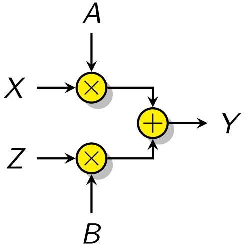

- In general, we have

$$Y = AX + BZ$$

for some transfer matrices $A \in R^{N\times n}$ and $B \in R^{N\times t}$.

Capacity of Matrix Channels over Finite Chain Rings

Our model is $Y = AX + BZ$.

- A well-defined discrete memoryless channel with input alphabet $R^{n\times\mu}$, output alphabet $R^{N\times\mu}$ and channel transition probability $p_{Y|X}$ is obtained once a joint distribution for $p_{Z,A,B|X}$ is specified.
- The capacity of this channel is given, as usual, by

$$C = \max_{p_X} I(X;Y)$$

where p_X is the input distribution. (We will take logarithms to base q, so the capacity is given in q-qary symbols per channel use.)

Asymptotic Capacity, $\overline{C}$

How does capacity scale with packet length?
Given a channel with a given n, N, μ, and t, we define the kth extension as the channel in which the transmitter sends kn packets of shape $k\mu$, the receiver gather kN packets of this shape, the noise matrix has kt rows, and the channel law is suitably generalized, giving capacity C_k.

Definition

We define the asymptotic capacity as

$$\overline{C} = \lim_{k\to\infty} \frac{1}{(kn)|k\mu|} C_k = \frac{1}{n|\mu|} \lim_{k\to\infty} \frac{C_k}{k^2}.$$

Note that $\overline{C}$ is normalized such that $\overline{C} = 1$ if the channel is noiseless (i.e., $A = I$ and $Z = 0$).

The Independent Transfer Model

Let τ be an s-shape such that $\tau \preceq t, \mu$.
We study the case where:

- the transfer matrix A is uniform over $\mathrm{GL}_n(R)$ (in particular, $N = n$),
- B is uniform over $\mathcal{T}_t(R^{n\times t})$,
- Z is uniform over $\mathcal{T}_\tau(R^{t\times\mu})$,
- X, A, B and Z are statistically independent.

In this case we can re-write the channel model as

$$Y = A(X + A^{-1}BZ) = A(X + W),$$

where $A \in \mathrm{GL}_n(R)$ and $W \triangleq A^{-1}BZ \in \mathcal{T}_\tau(R^{n\times\mu})$ are chosen uniformly at random and independently from any other variables.

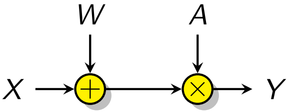

MMC: Model

First warmup problem

The multiplicative matrix channel (MMC):

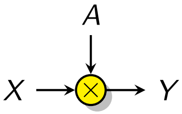

$$Y = AX$$

where

- $X, Y \in R^{n\times\mu}$;
- $A \sim \mathrm{Unif}[\mathrm{GL}_n(R)]$;
- A and X are independent.

MMC: Exact Capacity

It is easy to find the capacity of a channel defined by a group action.

- Let $\mathcal{G}$ be a finite group that acts on a finite set $\mathcal{S}$.
- Consider a channel with input $X \in \mathcal{S}$, output $Y \in \mathcal{S}$ and channel law $Y = AX$, where $A \sim \text{Unif}[\mathcal{G}]$ and A and X are independent.
- The capacity of this channel is

$$C = \log |\mathcal{S}/\mathcal{G}|,$$

  where $|\mathcal{S}/\mathcal{G}|$ is the number of orbits of the action.
- One capacity-achieving input distribution is to sample uniformly over a complete system of orbit-representatives.

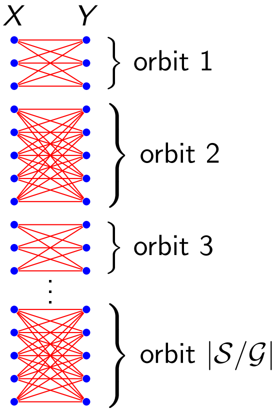

MMC: Exact Capacity

In the case of the MMC,

- $GL_n(R)$ acts on $R^{n\times\mu}$ by left-multiplication.
- The orbits are the sets of matrices that share the same row module.
- The number of such orbits is the number of submodules of R^μ with shape $\preceq n, \mu$.

Theorem

The capacity of the MMC, in q-ary symbols per channel use, is given by

$$C_{MMC} = \log_q \sum_{\lambda \preceq n,\mu} \begin{bmatrix} \mu \\ \lambda \end{bmatrix}_q .$$

A capacity-achieving code $\mathcal{C} \subseteq R^{n\times\mu}$ consists of all possible row canonical forms in $R^{n\times\mu}$.

(This scheme encodes information in the choice of submodules, generalizing the "transmission via subspaces" approach of [KK08].)

MMC: Asymptotic Capacity

The capacity C_{MMC} is bounded by

$$\sum_{i=1}^{s} \kappa_i(\mu_i - \kappa_i) \le C_{\mathrm{MMC}} \le \sum_{i=1}^{s} \kappa_i(\mu_i - \kappa_i) + \log_q 4^s \binom{n+s}{s} \quad (2)$$

where $\kappa_i = \min\{n, \lfloor \mu_i/2 \rfloor\}$.

Theorem

$$\overline{C}_{MMC} = \frac{\sum_{i=1}^{s} \kappa_i(\mu_i - \kappa_i)}{n|\mu|},$$

where $\kappa_i = \min\{n, \lfloor \mu_i/2 \rfloor\}$.

The choice of subshape κ essentially maximizes the number of principal row canonical forms having fixed subshape.
Thus, asymptotically, capacity can be achieved by always transmitting principal row canonical forms with a fixed subshape!

MMC: Encoding and Decoding

Let $\kappa = (\kappa_1, \ldots, \kappa_s)$ with $\kappa_i = \min\{n, \lfloor \mu_i/2 \rfloor\}$.

- Encoding: choose the input matrix X from the set of principal RCFS for $\mathcal{T}_\kappa(R^{n\times\mu})$ using the π-adic decomposition given earlier. The encoding rate is

$$R_{\mathrm{MMC}} = \sum_{i=1}^{s} \kappa_i(\mu_i - \kappa_i).$$

- Decoding: upon receiving $Y = AX$, the decoder simply computes the row canonical form of Y. The decoding is always correct by the uniqueness of the row canonical form.

This coding scheme achieves the asymptotic capacity $\overline{C}_{\mathrm{MMC}}$.

AMC: Model

Second warmup problem

The additive matrix channel (AMC):

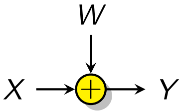

$$Y = X + W$$

where

- $X, Y \in R^{n \times \mu}$;
- $W \sim \text{Unif}[\mathcal{T}_\tau(R^{n \times \mu})]$;
- W and X are independent.

AMC: Exact Capacity

The AMC is an example of a discrete symmetric channel.

Theorem

The capacity of the AMC, in q-ary symbols per channel use, is given by

$$C_{AMC} = \log_q |R^{n\times\mu}| - \log_q |\mathcal{T}_\tau(R^{n\times\mu})|,$$

achieved by the uniform input distribution.

AMC: Asymptotic Capacity

The capacity C_{AMC} is bounded by

$$\sum_{i=1}^{s}(n-\tau_i)(\mu_i-\tau_i) - \log_q 4^s \prod_{i=0}^{\tau_s-1}(1-q^{i-n}) < C_{\mathsf{AMC}} < \sum_{i=1}^{s}(n-\tau_i)(\mu_i-\tau_i) - \log_q \prod_{i=0}^{\tau_s-1}(1-q^{i-n}).$$

Theorem

The asymptotic capacity $\overline{C}_{AMC}$ is given by

$$\overline{C}_{AMC} = \frac{\sum_{i=1}^{s}(n-\tau_i)(\mu_i-\tau_i)}{n|\mu|}.$$

AMC: Error-trapping Encoding

We focus on the special case when $\tau = t = (t, \ldots, t)$. Set $v \geq t$ and transmit a matrix X of the form

$$X = \begin{bmatrix} 0 & 0 \\ 0 & U_{(n-v)\times(m-v)} \end{bmatrix}.$$

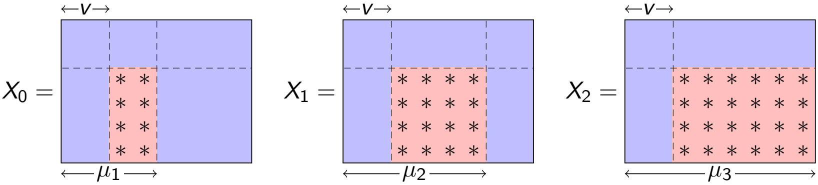

Clearly

$$R_{\text{AMC}} = \sum_{i=1}^{s} (n - v)(\mu_i - v).$$

AMC: Error-trapping Decoding

Write

$$W = \begin{bmatrix} W_1 & W_2 \\ W_3 & W_4 \end{bmatrix}.$$

Suppose shape $W_1 = t$. Then, since shape $W = t$ also, the pivots of W are entirely contained in W_1. Since the row canonical form of W has t nonzero rows, this means that the upper rows of W can cancel the lower rows, i.e., for some matrix V we have

$$\begin{bmatrix} I & 0 \\ V & I \end{bmatrix} \begin{bmatrix} W_1 & W_2 \\ W_3 & W_4 \end{bmatrix} = \begin{bmatrix} W_1 & W_2 \\ 0 & 0 \end{bmatrix}$$

Indeed, V can be chosen so that $VW_1 = -W_3$, which automatically forces $VW_2 = -W_4$ (since if $VW_2 + W_4 \neq 0$, W would have pivots outside of W_1).

Applying this transformation to $Y = X + W$ yields

$$\begin{bmatrix} I & 0 \\ V & I \end{bmatrix} \begin{bmatrix} W_1 & W_2 \\ W_3 & U + W_4 \end{bmatrix} = \begin{bmatrix} W_1 & W_2 \\ 0 & U \end{bmatrix},$$

exposing the user's data matrix U.

AMC: Error-trapping Decoding (cont'd)

In summary:

- The decoder observes W_1, W_2, and W_3 thanks to the error traps.
- If shape $W_1 = t$, then the decoder applies the transformation on the previous slide to expose U.
- If shape $W_1 \neq t$, a decoding failure (detected error) is declared.

The probability of decoding failure $P_f = P[\text{shape}\, W_1 \neq t]$ is bounded as

$$P_f < \frac{2t}{q^{1+v-t}}.$$

If we set v such that $v - t \to \infty$, and $\frac{v-t}{m} \to 0$, as $m \to \infty$, then we have $P_f \to 0$ and $\bar{R}_{\text{AMC}} = \frac{R_{\text{AMC}}}{n|\mu|} \to \bar{C}_{\text{AMC}}$.

Theorem

This coding scheme can achieve the asymptotic capacity of the AMC for the special case when $\tau = t$.

AMMC: Model

Now to the main event:

The additive-multiplicative matrix channel (AMMC):

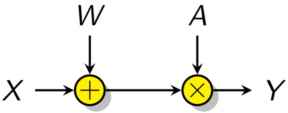

$$Y = A(X + W)$$

where

- $X, Y \in R^{n \times \mu}$;
- $W \sim \mathrm{Unif}[\mathcal{T}_\tau(R^{n \times \mu})]$;
- $A \sim \mathrm{Unif}[\mathrm{GL}_n(R)]$;
- A, X and W are independent.

Remark: This model is statistically identical to $Y = AX + Z$, where $Z \sim \mathrm{Unif}[\mathcal{T}_\tau(R^{n \times \mu})]$

AMMC: Upper Bound on Capacity

Theorem

The capacity of the AMMC, in q-ary symbols per channel use, is upper-bounded by

$$C_{AMMC} \leq \sum_{i=1}^{s} (\mu_i - \xi_i)\xi_i + \sum_{i=1}^{s} (n - \mu_i)\tau_i + 2s \log_q 4 + \log_q \binom{n+s}{s}$$

$$+ \log_q \binom{\tau_s + s}{s} - \log_q \prod_{i=0}^{\tau_s - 1} (1 - q^{i-n}), \ \textit{where } \xi_i = \min\{n, \lfloor \mu_i/2 \rfloor\}.$$

In particular, when $\mu \succeq 2n$, the upper bound reduces to

$$C_{\mathrm{AMMC}} \leq \sum_{i=1}^{s} (n - \tau_i)(\mu_i - n) + 2s \log_q 4$$

$$+ \log_q \binom{n+s}{s} + \log_q \binom{\tau_s + s}{s} - \log_q \prod_{i=0}^{\tau_s - 1} (1 - q^{i-n}).$$

AMMC: Asymptotic Capacity

Theorem

When $\mu \succeq 2n$, the asymptotic capacity $\overline{C}_{AMMC}$ is upper-bounded by

$$\overline{C}_{AMMC} \leq \frac{\sum_{i=1}^{s}(n-\tau_i)(\mu_i - n)}{n|\mu|}.$$

AMMC: Coding Scheme

We again focus on the special case when $\tau = t$, and combine the two strategies for the MMC and the AMC.

To **encode**, construct X as

$$X = \begin{bmatrix} 0 & 0 \\ 0 & \bar{X} \end{bmatrix},$$

where $\bar{X}$ is chosen from the set of principal row canonical forms for $\mathcal{T}_\kappa(R^{(n-v)\times(\mu-v)})$ by the previous construction.

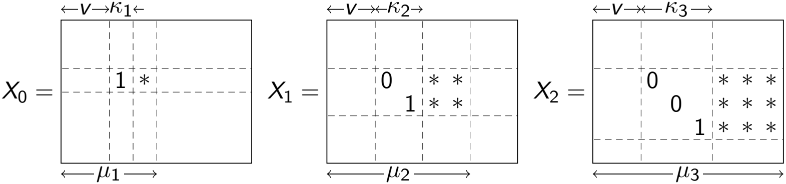

We have $R_{\text{AMMC}} = \sum_{i=1}^{s} \kappa_i(\mu_i - v - \kappa_i)$. In particular, when $\mu \succeq 2n$, we have $\lfloor(\mu_i - v)/2\rfloor \geq n - v$ for all i. Thus, $\kappa_i = n - v$ for all i, and the encoding rate is $R_{\text{AMMC}} = \sum_{i=1}^{s}(n-v)(\mu_i - n)$.

AMMC: Coding Scheme (cont'd)

To **decode**, we must recover $\bar{X}$ from $Y = A(X + W)$.
If we had $X + W$, we could use the error-trapping decoder to recover

$$\begin{bmatrix} W_1 & W_2 \\ 0 & \bar{X} \end{bmatrix}.$$

But we have Y, not $X + W$. However, since A is invertible, $\mathrm{RCF}(Y) = \mathrm{RCF}(X + W)$, and one easily sees that

$$\mathrm{RCF}(X + W) = \begin{bmatrix} \bar{W}_1 & \bar{W}_2 \\ 0 & \bar{X} \\ 0 & 0 \end{bmatrix},$$

where the bottom $v - t$ rows are all zero.
In summary:

- The decoder first computes $\mathrm{RCF}(Y)$.
- It then checks the condition shape $W_1 = t$.
- If the condition does not hold, a decoding failure is declared, otherwise the decoder outputs $\bar{X}$.

Conclusions

- Nested-lattice-based physical layer network coding naturally transforms wireless multiple-access channels with random fading into random linear network coding channels.
- The algebraic structure of Λ/Λ' is that of a module over a ring.
- In many cases, the ring is a finite-chain ring, so end-to-end error control (for random errors) can be handled using a matrix-channel approach, with simple and asymptotically efficient coding schemes.

Open Problems

- Relaxing the assumption on A
  - What if A is not invertible?
- Relaxing the assumption on W
  - What if W has shape other than $\tau = t$.?
- Adversarial error models
  - Always correcting errors when shape$(W) \leq \tau$?
- Rank-metric codes over finite chain rings
  - Which properties can be preserved?

Backup Slides

Physical-Layer Network Coding

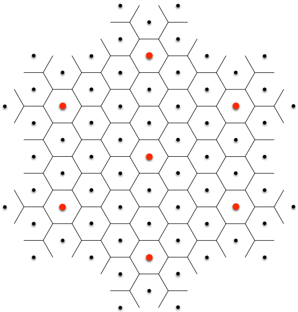

Motivation

Current Wireless

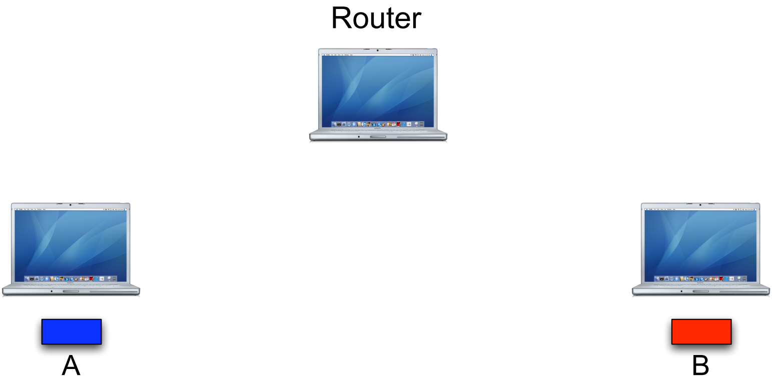

Current Wireless

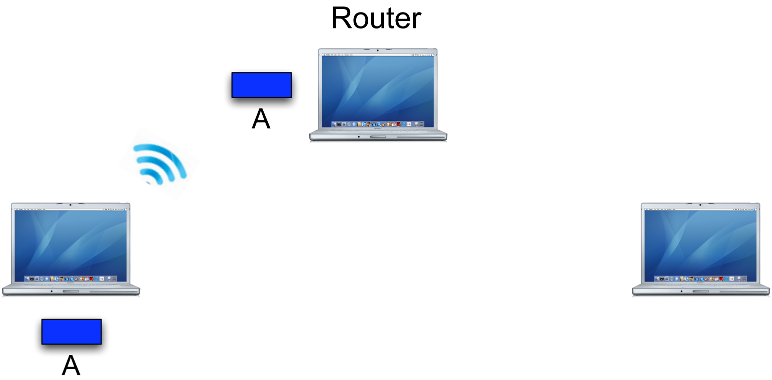

Current Wireless

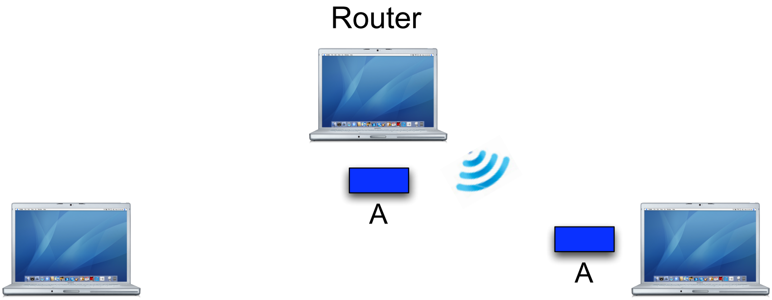

Current Wireless

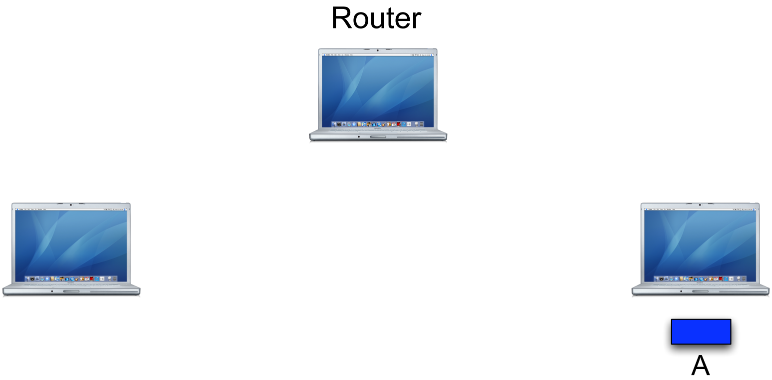

Current Wireless

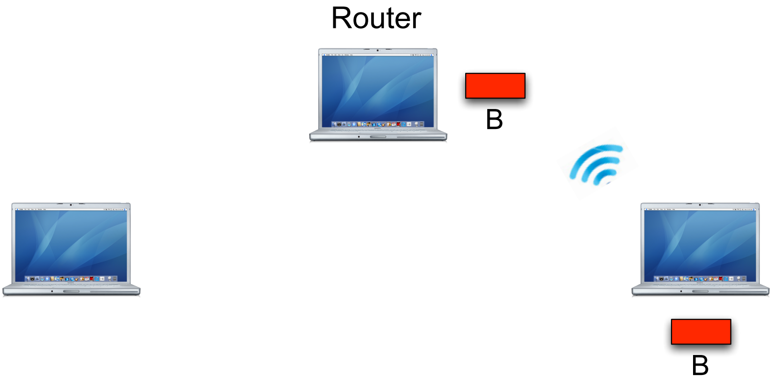

Current Wireless

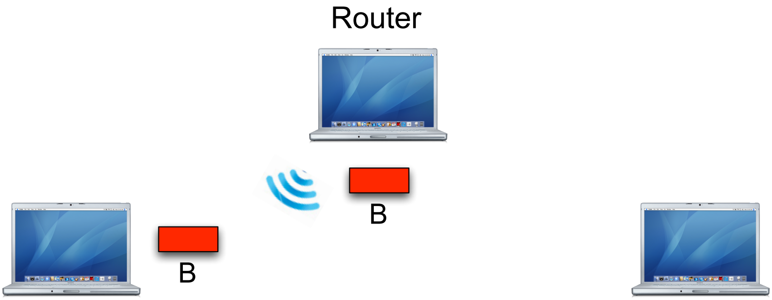

Current Wireless

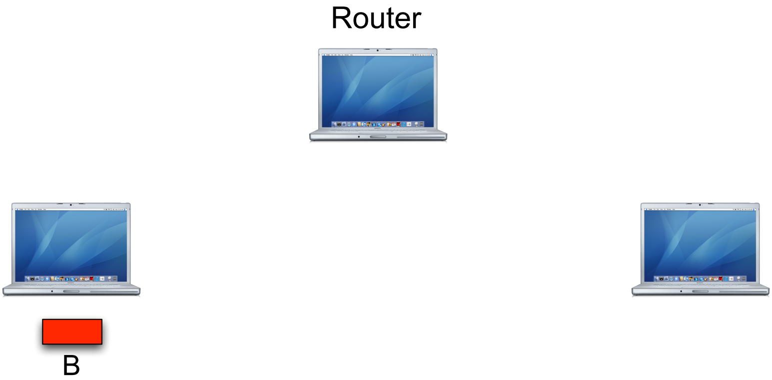

Current Wireless

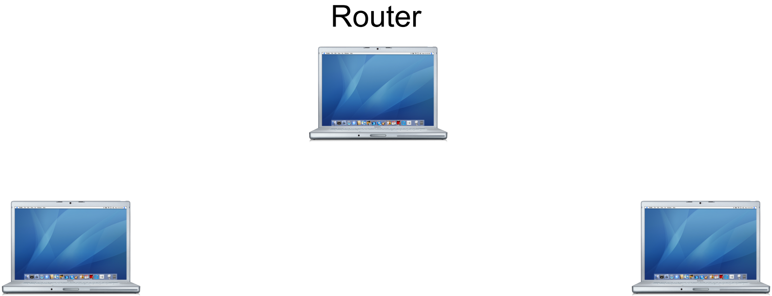

Routing requires 4 time slots

Network Coding

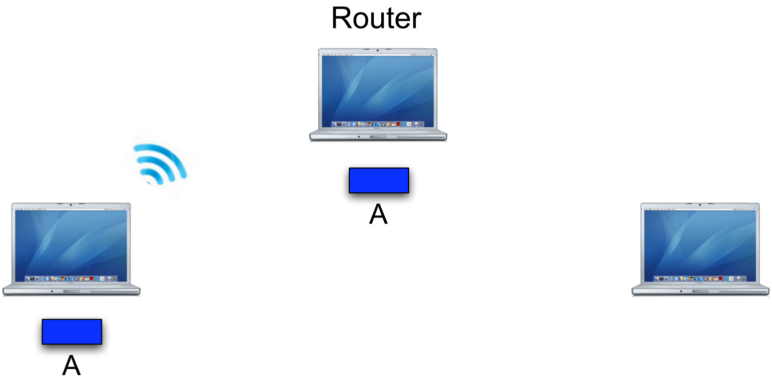

Network Coding

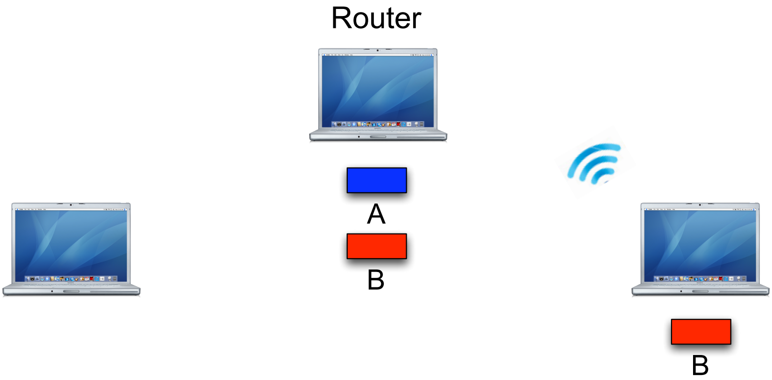

Network Coding

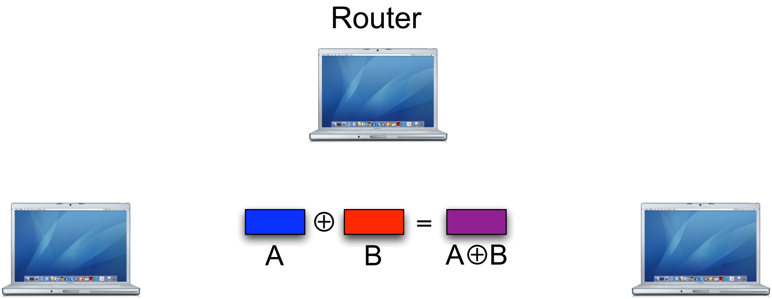

Network Coding

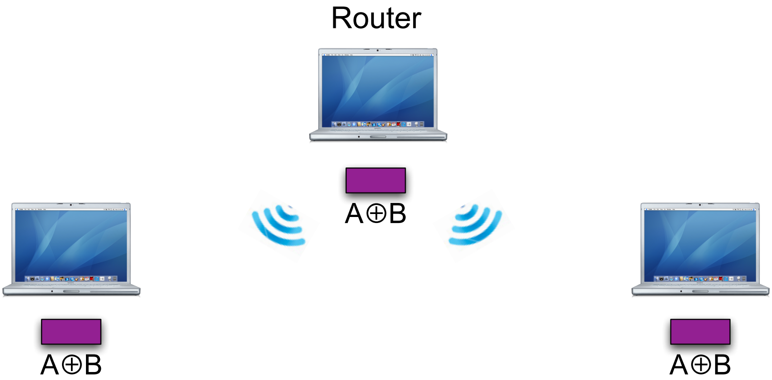

Network Coding

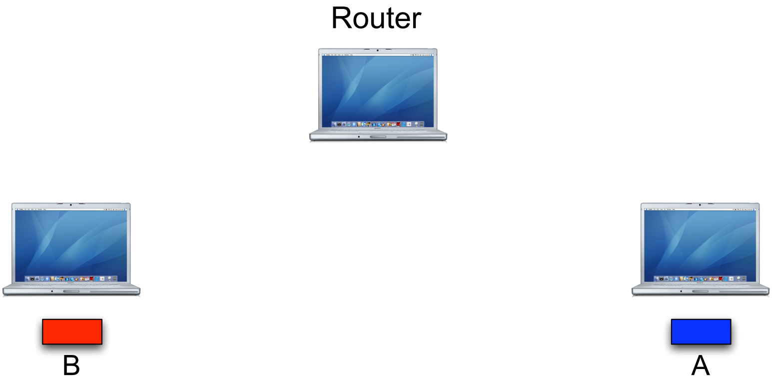

Network Coding

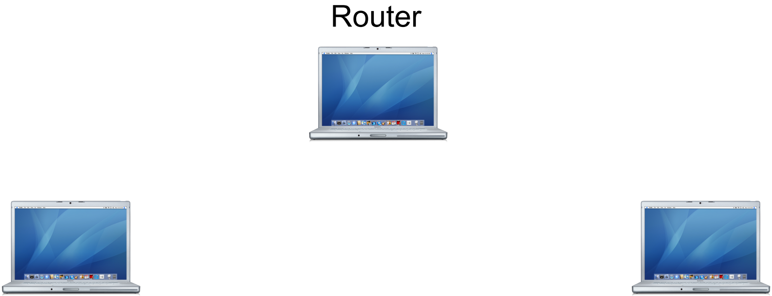

Network coding requires 3 time slots

Network Coding

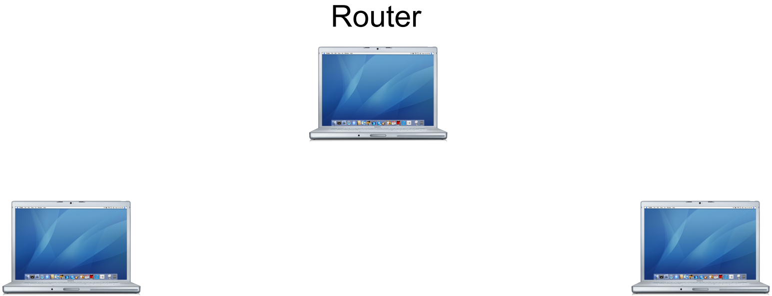

Network coding requires 3 time slots. Can we do better?

Physical-Layer Network Coding

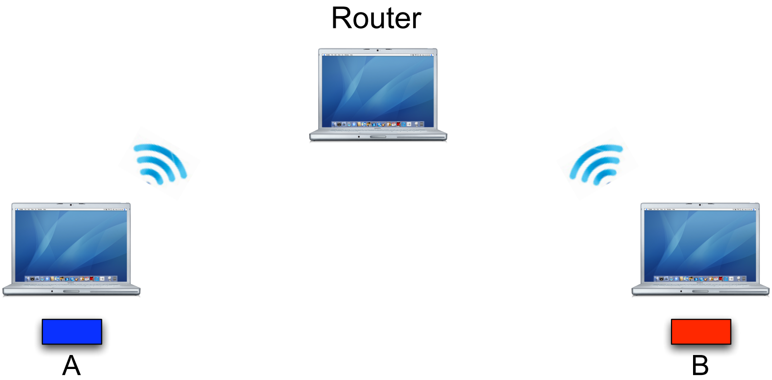

Physical-Layer Network Coding

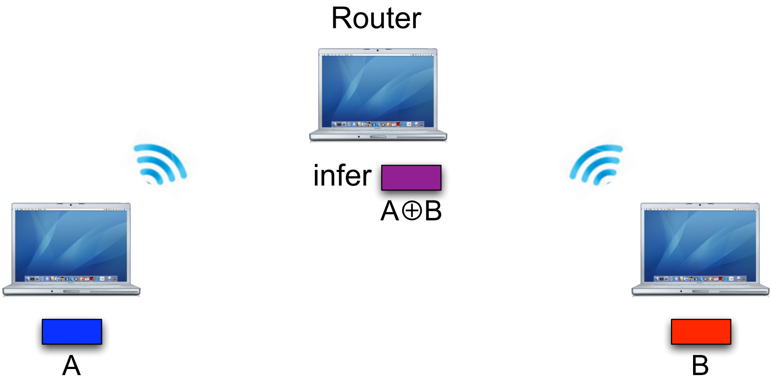

Physical-Layer Network Coding

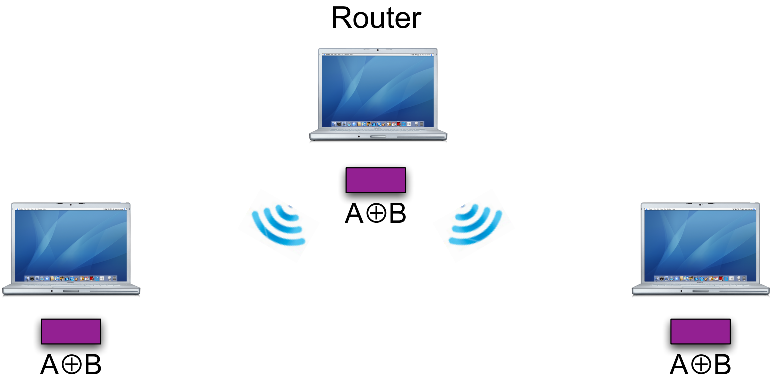

Physical-Layer Network Coding

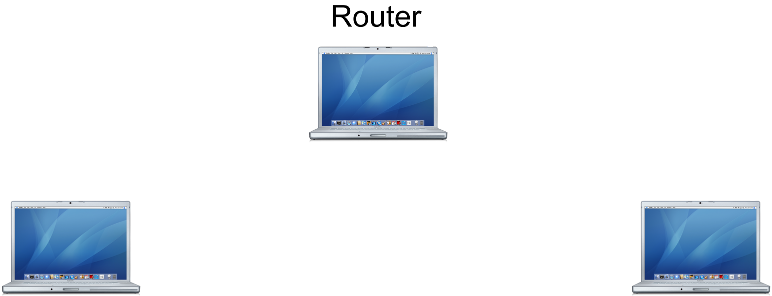

Physical-layer network coding requires 2 time slots

It Is More Than Going From 3 to 2

- **A new way of dealing with interference**

  process interference instead of avoiding it

- **Can be extended to large networks**

  each relay infers some linear combination

Extension to Large Wireless Networks

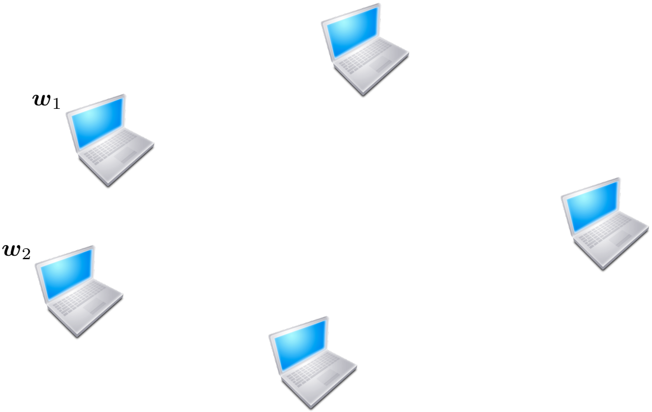

Extension to Large Wireless Networks

Extension to Large Wireless Networks

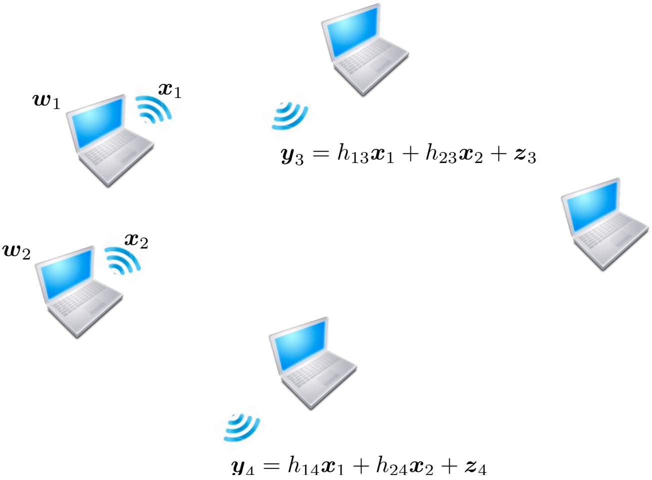

Extension to Large Wireless Networks

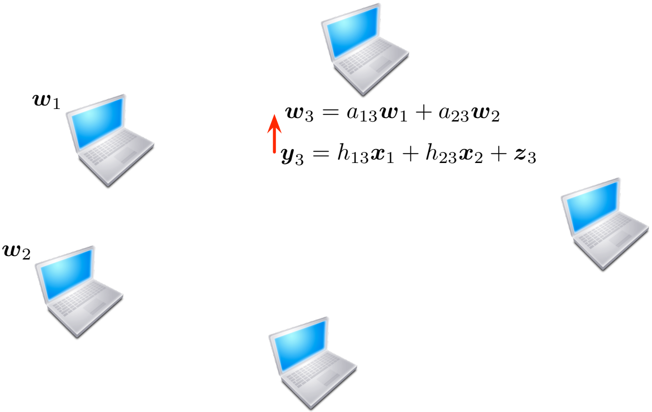

$$\boldsymbol{w}_3 = a_{13}\boldsymbol{w}_1 + a_{23}\boldsymbol{w}_2$$

$$\boldsymbol{y}_3 = h_{13}\boldsymbol{x}_1 + h_{23}\boldsymbol{x}_2 + \boldsymbol{z}_3$$

$$\boldsymbol{w}_4 = a_{14}\boldsymbol{w}_1 + a_{24}\boldsymbol{w}_2$$

$$\boldsymbol{y}_4 = h_{14}\boldsymbol{x}_1 + h_{24}\boldsymbol{x}_2 + \boldsymbol{z}_4$$

Extension to Large Wireless Networks

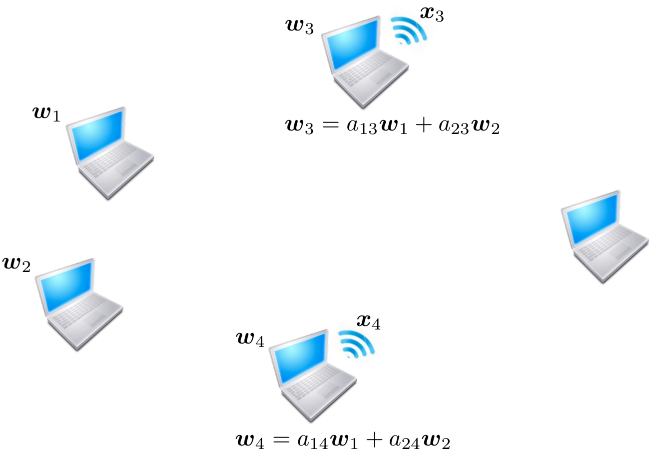

Extension to Large Wireless Networks

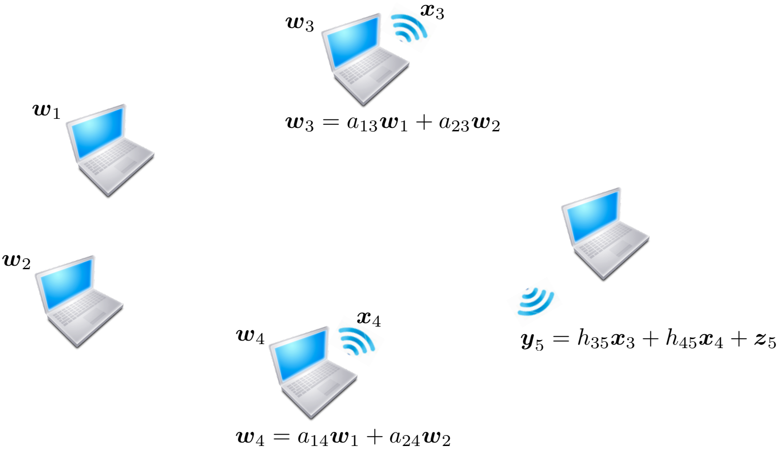

Extension to Large Wireless Networks

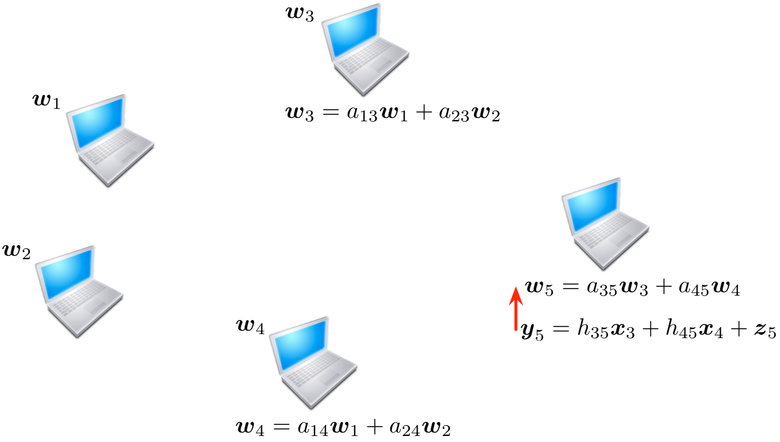

Part 1: First PNC schemes

Example 1 (BPSK, $h_1 = h_2 = 1$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

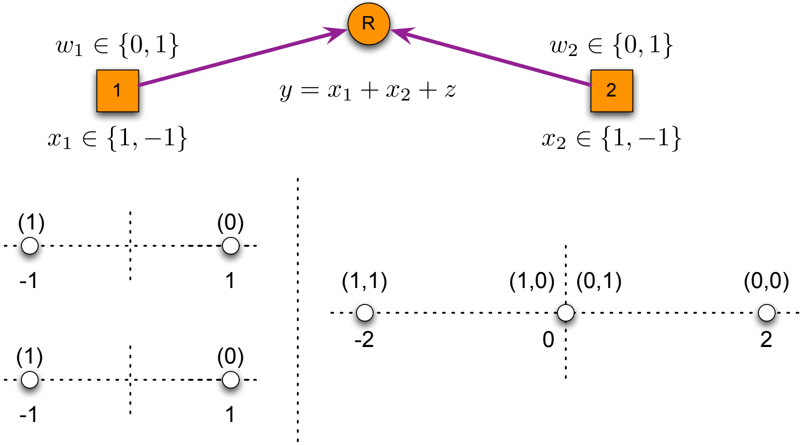

Example 1 (BPSK, $h_1 = h_2 = 1$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

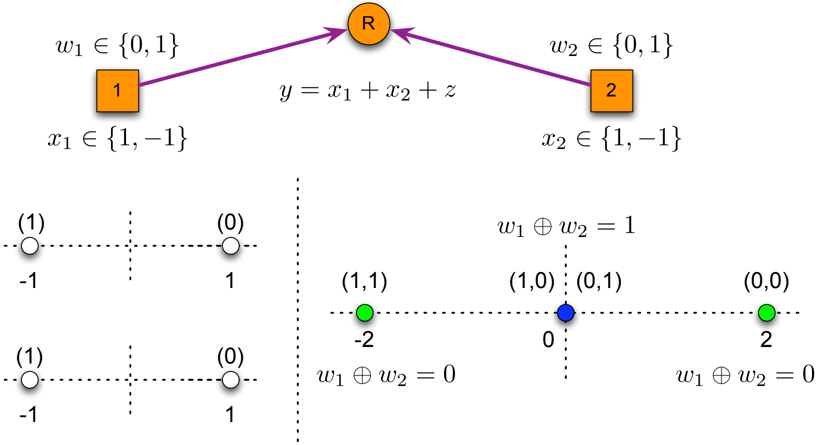

Example 1 (BPSK, $h_1 = h_2 = 1$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

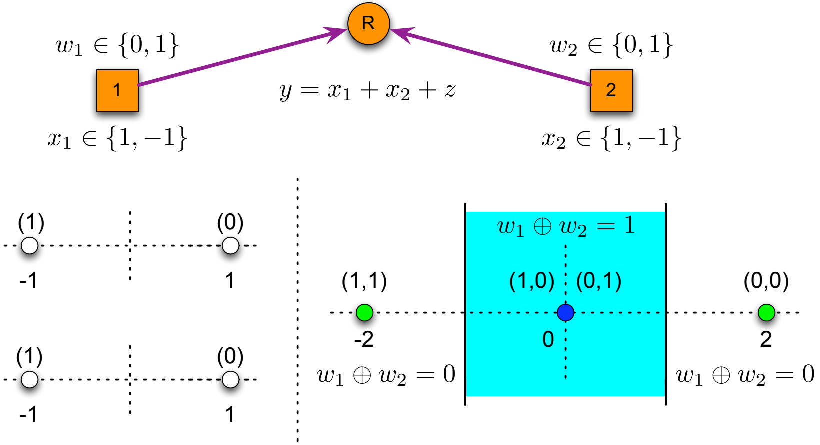

Example 2 (QPSK, $h_1 \approx h_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

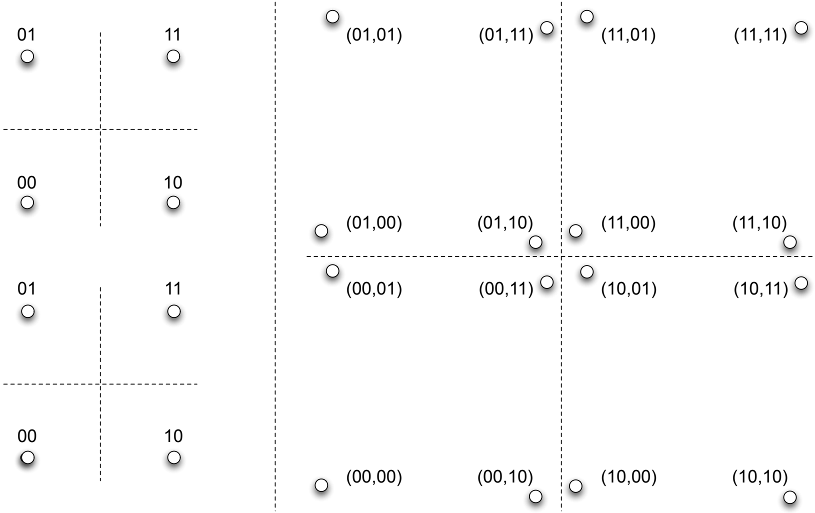

Example 2 (QPSK, $h_1 \approx h_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

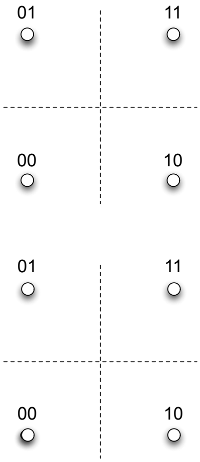

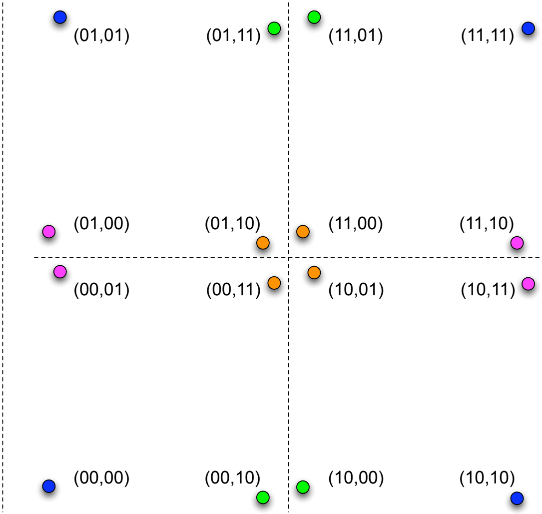

Example 2 (QPSK, $h_1 \approx h_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

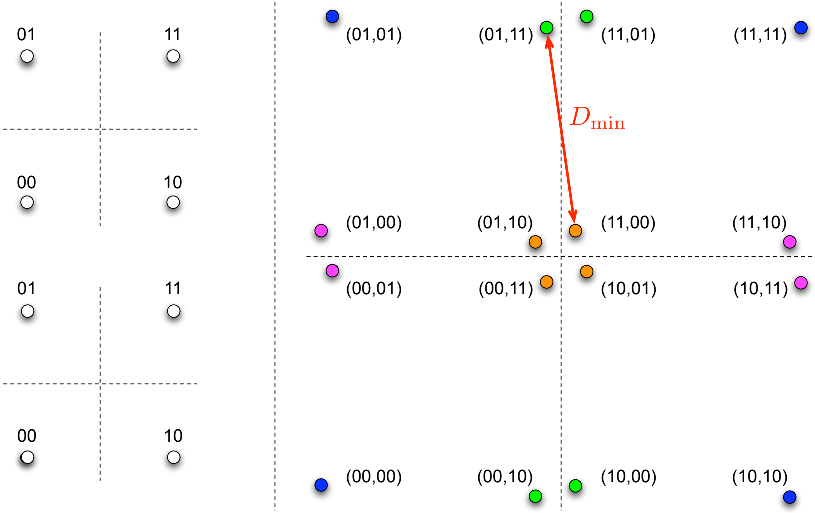

Example 3 (QPSK, $h_1 \approx ih_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

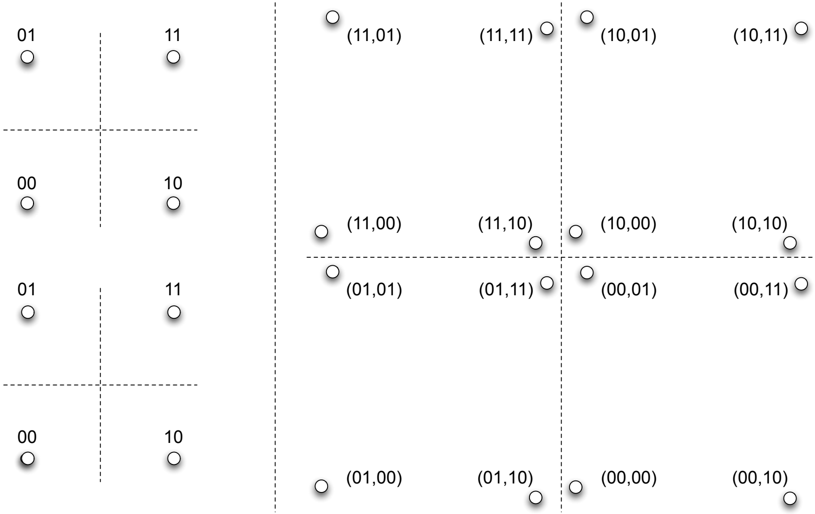

Example 3 (QPSK, $h_1 \approx ih_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

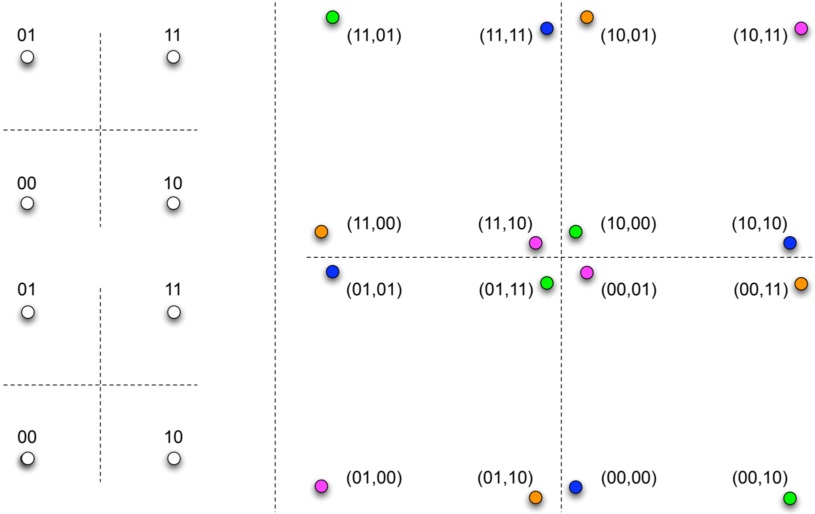

Example 3 (QPSK, $h_1 \approx ih_2$)

Zhang-Liew-Lam 2006, Popovski-Yomo 2006

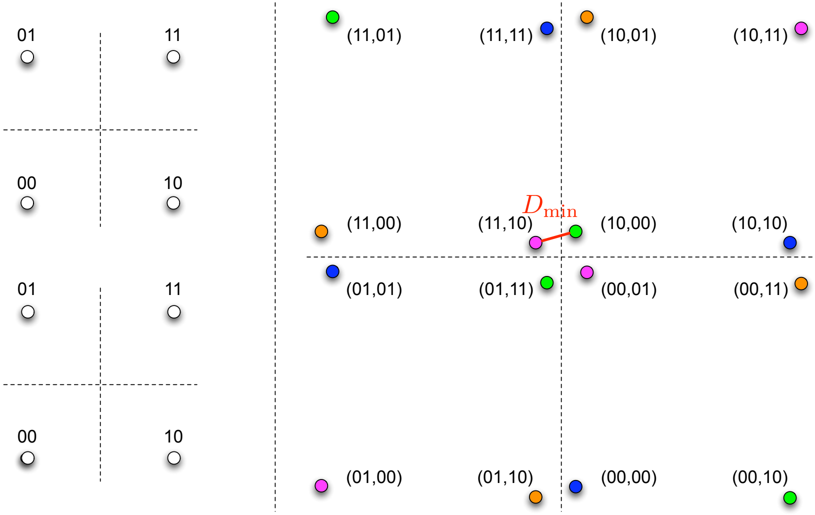

Limitation of Original PNC Schemes

Limitation: phase misalignment $\Rightarrow$ bad performance

Solution 1: require phase synchronization

Solution 2: mitigate phase misalignment by moving beyond XOR

- Popovski & Yomo 2007
- Koike-Akino-Popovski-Tarokh 2008

Solution 3: mitigate phase misalignment by compute-and-forward

- Nazer & Gastpar 2007

Example 4 (QPSK, $h_1 \approx ih_2$)

Koike-Akino-Popovski-Tarokh: $(ab, cd) \to ab \oplus dc$

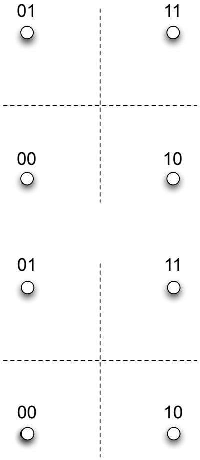

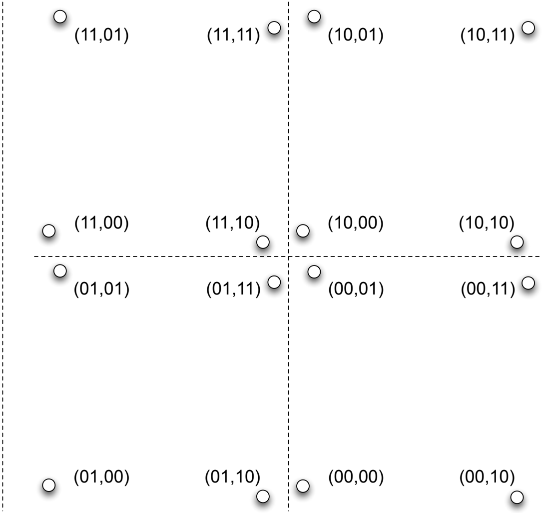

Example 4 (QPSK, $h_1 \approx ih_2$)

Koike-Akino-Popovski-Tarokh: $(ab, cd) \to ab \oplus dc$

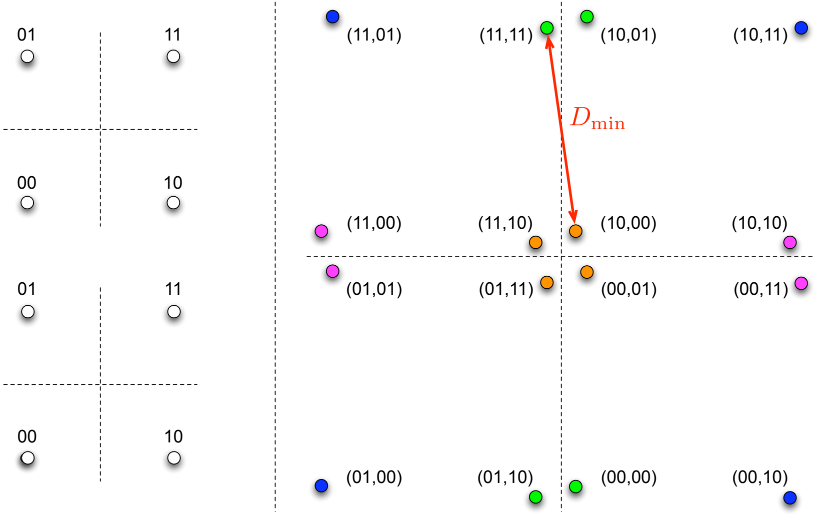

Part 2: Compute-and-Forward

Compute-and-Forward Relaying Strategy

Nazer & Gastpar's Approach (2006)

- Voronoi constellations based on Erez-Zamir's construction
- Main result: achievable rates for one-hop networks
- CSI only at the receivers but not at the transmitters

Similar Approaches

- Narayanan-Wilson-Sprintson (2007)
- Nam-Chung-Lee (2008)
- Wilson-Narayanan (2009)

Remark: all of these are based on Erez-Zamir's construction of Voronoi constellations

Voronoi Constellations in One Slide

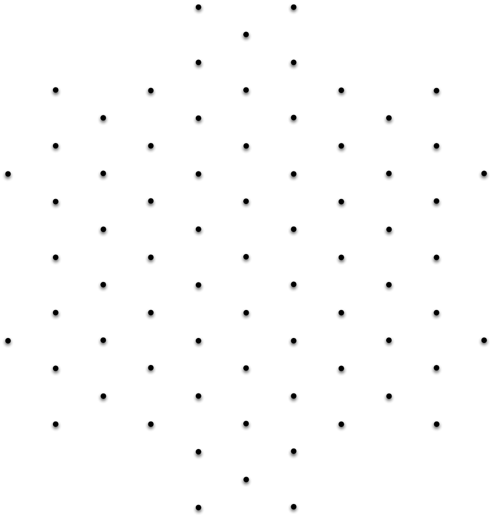

pick a fine lattice Λ

Voronoi Constellations in One Slide

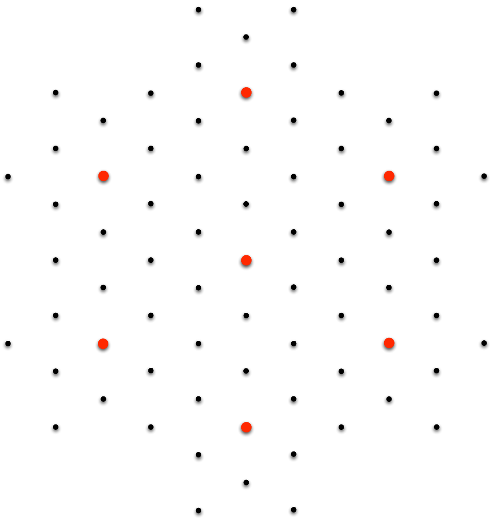

pick a coarse lattice Λ'

Voronoi Constellations in One Slide

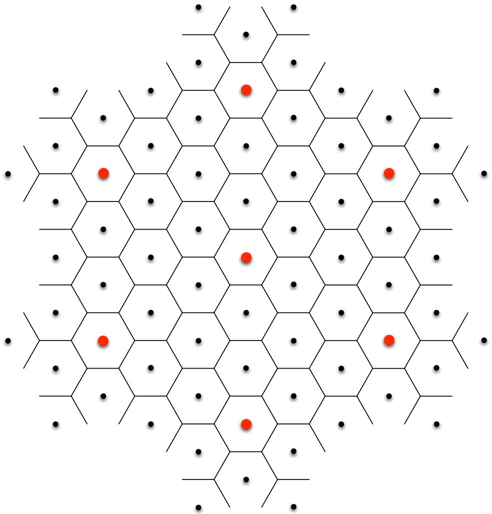

Voronoi region for the fine lattice Λ

Voronoi Constellations in One Slide

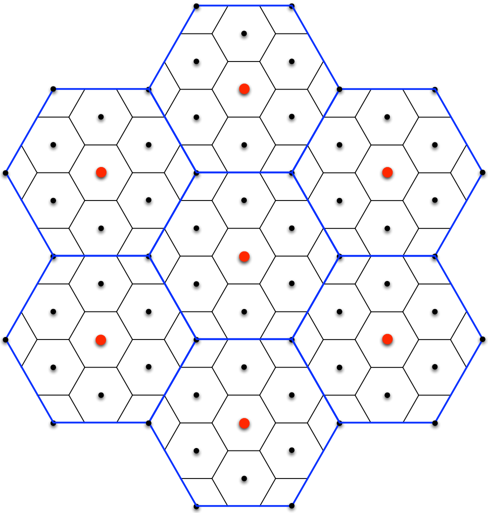

Voronoi region for the coarse lattice Λ'

Key Idea: the Case of Integer Channel Gains

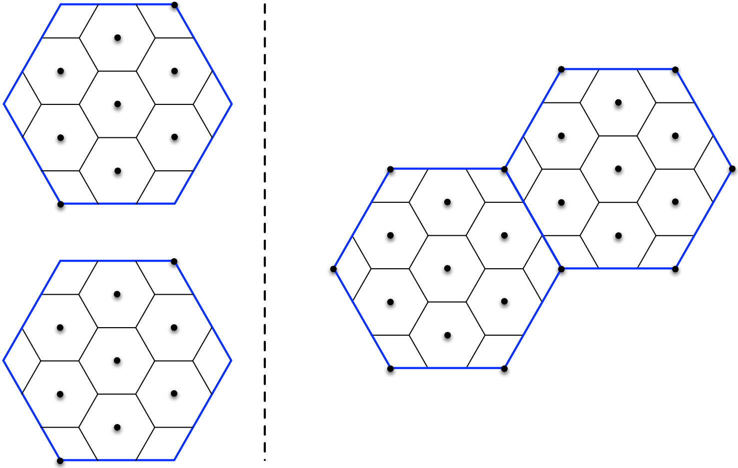

Each transmitter applies the same Voronoi constellation Λ/Λ'

Key Idea: the Case of Integer Channel Gains

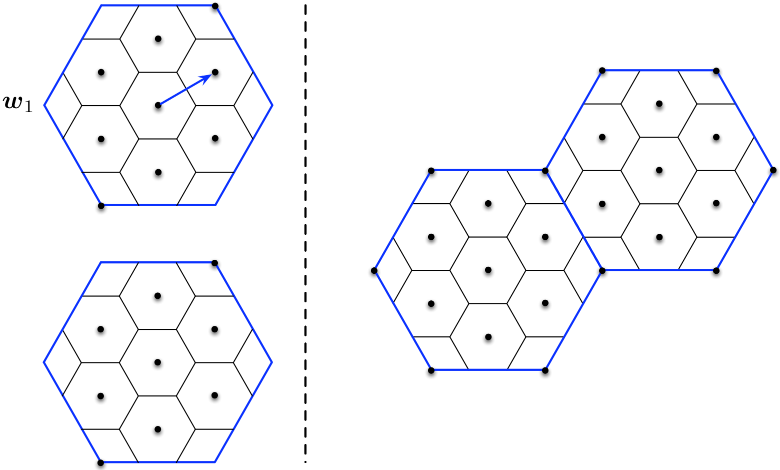

Transmitter 1 maps $\mathbf{w}_1$ to a constellation point

Key Idea: the Case of Integer Channel Gains

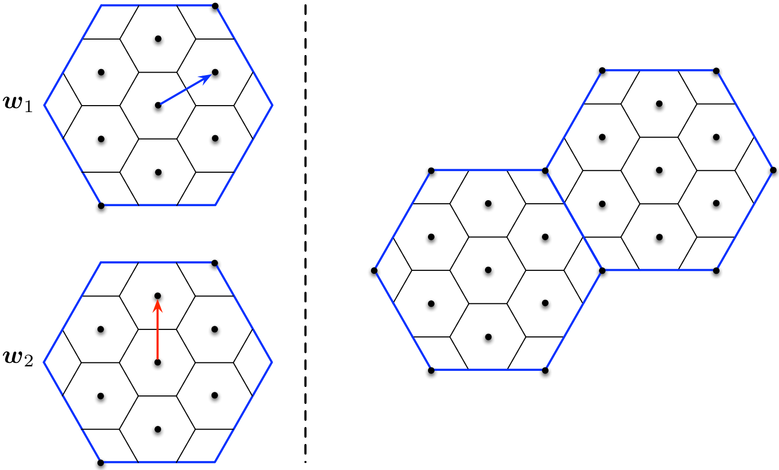

Transmitter 2 maps $\mathbf{w}_2$ to a constellation point

Key Idea: the Case of Integer Channel Gains

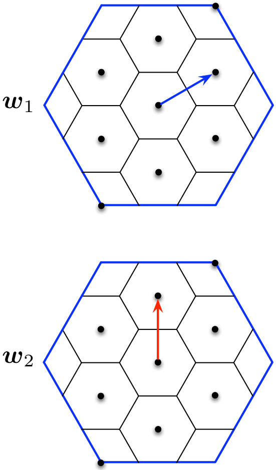

$$\boldsymbol{y} = h_1 \boldsymbol{x}_1 + h_2 \boldsymbol{x}_2 + \boldsymbol{z}$$

$$(h_1, h_2) = (2, 1)$$

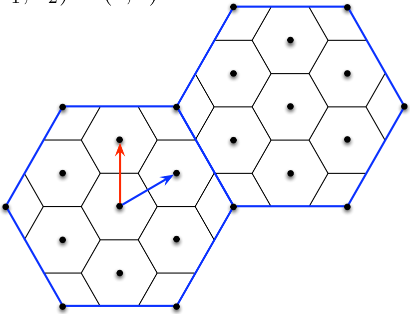

The channel is given by $\mathbf{y} = 2\mathbf{x}_1 + \mathbf{x}_2 + \mathbf{z}$

Key Idea: the Case of Integer Channel Gains

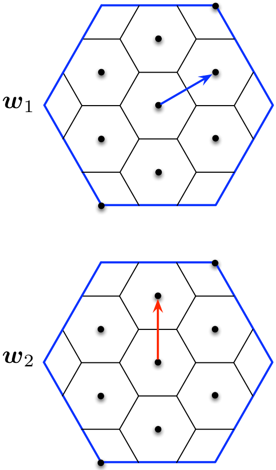

$$\boldsymbol{y} = h_1\boldsymbol{x}_1 + h_2\boldsymbol{x}_2 + \boldsymbol{z}$$

$$(h_1, h_2) = (2, 1)$$

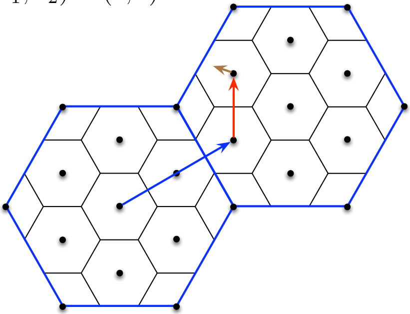

Hence, the received signal **y** is like this

Key Idea: the Case of Integer Channel Gains

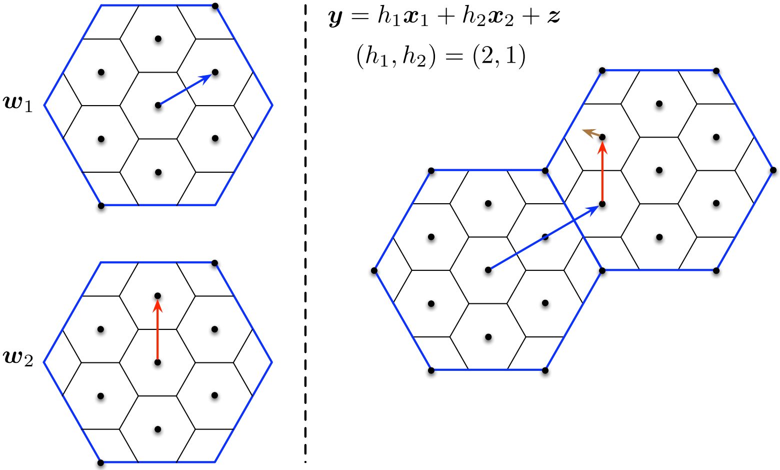

$$\boldsymbol{y} = h_1 \boldsymbol{x}_1 + h_2 \boldsymbol{x}_2 + \boldsymbol{z}$$

$$(h_1, h_2) = (2, 1)$$

But how can we extract some information from **y**?

Key Idea: the Case of Integer Channel Gains

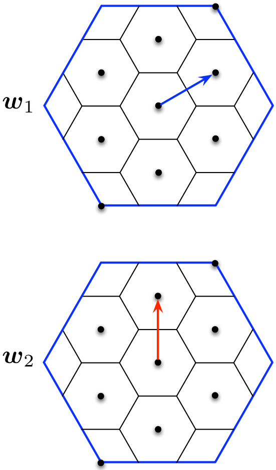

$$\boldsymbol{y} = h_1\boldsymbol{x}_1 + h_2\boldsymbol{x}_2 + \boldsymbol{z}$$

$$(h_1, h_2) = (2, 1)$$

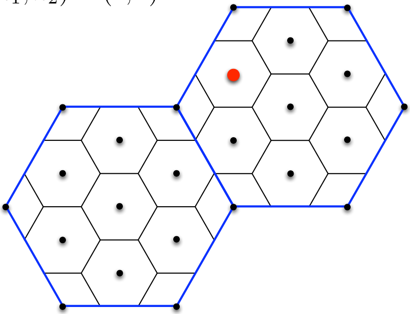

The receiver can decode $2\mathbf{x}_1 + \mathbf{x}_2$, if the noise $\mathbf{z}$ is small

Key Idea: the Case of Integer Channel Gains

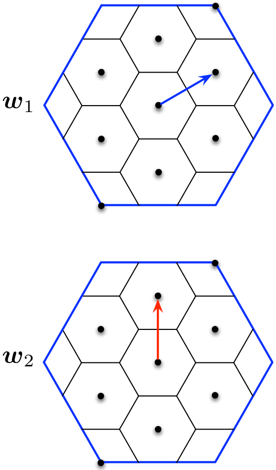

$$\boldsymbol{y} = h_1\boldsymbol{x}_1 + h_2\boldsymbol{x}_2 + \boldsymbol{z}$$

$$(h_1, h_2) = (2, 1)$$

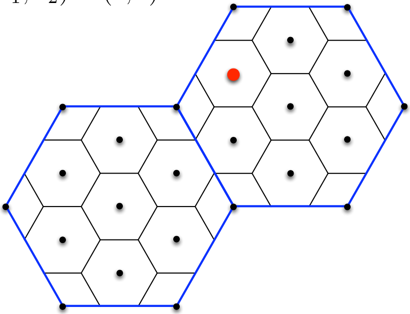

But we need a linear combination of messages...

Key Idea: the Case of Integer Channel Gains

Note that...

one can construct a one-to-one linear map between $\mathbb{F}_3^2$ and Λ/Λ'

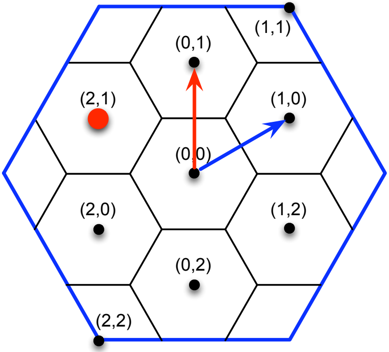

$$\boldsymbol{w}_1 = (1,0)$$

$$\boldsymbol{w}_2 = (0,1)$$

$$\hat{\boldsymbol{u}} = (2,1)$$

$$\hat{\boldsymbol{u}} = 2\boldsymbol{w}_1 + \boldsymbol{w}_2$$

Hence, an integer combination of lattice points
$=$ a linear combination of messages

Key Idea: the Case of Real Channel Gains

What if channel gains are real numbers?

Applying a scaling operation $g(\mathbf{y}) = \alpha\mathbf{y}$

$$\begin{aligned}
\alpha\mathbf{y} &= \sum_{\ell} \alpha h_\ell \mathbf{x}_\ell + \alpha\mathbf{z} \\
&= \sum_{\ell} a_\ell \mathbf{x}_\ell + \underbrace{\sum_{\ell}(\alpha h_\ell - a_\ell)\mathbf{x}_\ell + \alpha\mathbf{z}}_{\text{effective noise}} \\
&= \sum_{\ell} a_\ell \mathbf{x}_\ell + \mathbf{n},
\end{aligned}$$

where $\{a_\ell\}$ are integers, and $\alpha \in \mathbb{R}$ is the scalar.
Thus, real-valued channel gains $\Rightarrow$ integer channel gains

Key Idea: the Case of Real Channel Gains

What if channel gains are real numbers?

Applying a scaling operation $g(\mathbf{y}) = \alpha\mathbf{y}$

$$\begin{aligned}\alpha\mathbf{y} &= \sum_{\ell} \alpha h_{\ell}\mathbf{x}_{\ell} + \alpha\mathbf{z} \\ &= \sum_{\ell} a_{\ell}\mathbf{x}_{\ell} + \underbrace{\sum_{\ell}(\alpha h_{\ell} - a_{\ell})\mathbf{x}_{\ell} + \alpha\mathbf{z}}_{\text{effective noise}} \\ &= \sum_{\ell} a_{\ell}\mathbf{x}_{\ell} + \mathbf{n},\end{aligned}$$

where $\{a_{\ell}\}$ are integers, and $\alpha \in \mathbb{R}$ is the scalar.

Thus, real-valued channel gains $\Rightarrow$ integer channel gains

But how shall we choose the scalar α?

see Nazer-Gastpar (*IEEE Trans. Info. Theory*, 2011) for details

Key Idea: the Case of Complex Channel Gains

What if channel gains are complex numbers?

Answer: lift real lattices to Gaussian lattices

Gaussian integers: $\mathbb{Z}[i] = \{a + bi : a, b \in \mathbb{Z}\}$

Gaussian lattices: any Gaussian integer combination of lattice points is a lattice point

Then, apply a scaling operation $g(\mathbf{y}) = \alpha\mathbf{y}$

$$\begin{aligned} \alpha\mathbf{y} &= \sum_{\ell} \alpha h_\ell \mathbf{x}_\ell + \alpha\mathbf{z} \\ &= \sum_{\ell} a_\ell \mathbf{x}_\ell + \underbrace{\sum_{\ell} (\alpha h_\ell - a_\ell)\mathbf{x}_\ell + \alpha\mathbf{z}}_{\text{effective noise}} \\ &= \sum_{\ell} a_\ell \mathbf{x}_\ell + \mathbf{n}, \end{aligned}$$

where $\{a_\ell\}$ are Gaussian integers, and $\alpha \in \mathbb{C}$ is the scalar

Main Result: Computation Rate

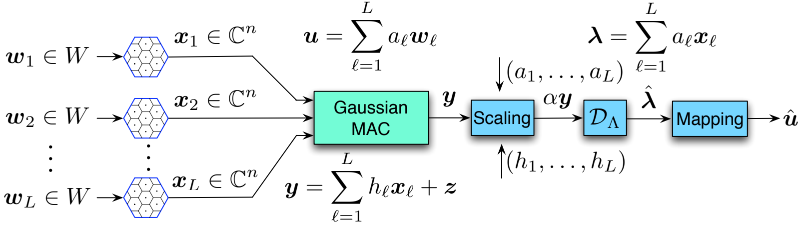

Computation Rate (Nazer-Gastpar)

$$R_{\text{comp}} = \log_2 \left(\frac{P}{P \sum_\ell \|\alpha h_\ell - a_\ell\|^2 + N_0 |\alpha|^2} \right)$$

Remark: Erez-Zamir's construction of Voronoi constellations $\Rightarrow$ asymptotically long block length and almost unbounded complexity

Research Problems

What if practical Voronoi constellations are used?

Goal: Practical Design for Compute-and-Forward

- Short block length and low complexity
- Example: wireless fading channel with short coherent time

Research Problems

What if practical Voronoi constellations are used?

Goal: Practical Design for Compute-and-Forward

- Short block length and low complexity
- Example: wireless fading channel with short coherent time

Related Work

- Ordentlich & Erez (2010)
- Hern & Narayanan (2011)
- Tunali & Narayanan (2011)
- Ordentlich-Zhan-Erez-Gastpar-Nazer (2011)
- Feng-Silva-Kschischang (2011)
- Emerging work includes Osmane & Belfiore (in submission)

Part 3: An Algebraic Approach

Algebraic Approach: Key Elements

R-Lattices

Let R be a discrete subring of $\mathbb{C}$ forming a principal ideal domain. Let $N \leq n$. An R-lattice of dimension N in $\mathbb{C}^n$ is defined as the set of all R-linear combinations of N linearly independent vectors, i.e.,

$$\Lambda = \{\mathbf{r}\mathbf{G}_\Lambda : \mathbf{r} \in R^N\},$$

where $\mathbf{G}_\Lambda \in \mathbb{C}^{N \times n}$ is called a generator matrix for Λ.

$R = \mathbb{Z}[\omega] \Rightarrow$ Eisenstein lattices; $R = \mathbb{Z}[i] \Rightarrow$ Gaussian lattices

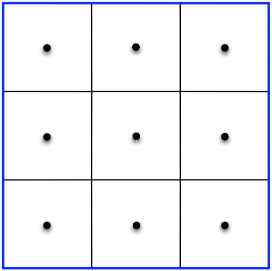

$$\mathbb{Z}[\omega] \triangleq \{a + b\omega : a, b \in \mathbb{Z}, \omega = e^{i2\pi/3}\}$$

$$\mathbb{Z}[i] \triangleq \{a + bi : a, b \in \mathbb{Z}\}$$

$$\Lambda = \mathbb{Z}[i]$$

$$\Lambda' = 3\mathbb{Z}[i]$$

Algebraic Approach: Key Concepts

Key Concepts

Message space W (with $|W| = |\Lambda/\Lambda'|$)

Labeling $\varphi : \Lambda \to W$ (consistent with Λ/Λ')

Embedding map $\bar{\varphi} : W \to \Lambda$ such that $\varphi(\bar{\varphi}(\mathbf{w})) = \mathbf{w}$

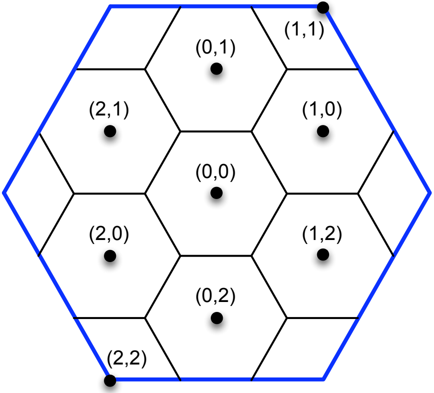

$$W = \mathbb{Z}_3 \times \mathbb{Z}_3$$

$$\varphi(\mathbf{w}\mathbf{G}_\Lambda) = \mathbf{w} \bmod 3$$

$$\bar{\varphi}(\mathbf{w}) = \mathbf{w}\mathbf{G}_\Lambda$$

Algebraic Approach: Key Property

Key Property

If the message space W is chosen carefully, then the labelling φ can be made linear.

In general, W can be chosen as $R/(\pi_1) \times \cdots \times R/(\pi_k)$, where $\pi_1, \ldots, \pi_k$ are invariant factors of Λ/Λ'.

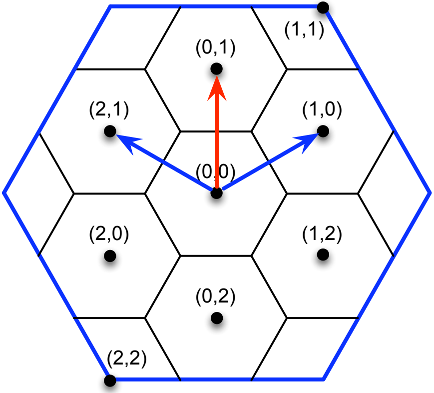

$$W = \mathbb{Z}_3 \times \mathbb{Z}_3$$

$$\varphi(\mathbf{w}\mathbf{G}_\Lambda) = \mathbf{w} \bmod 3$$

$$\varphi\Big(\bar{\varphi}(2,1) + \bar{\varphi}(1,0)\Big)$$
$$= (0,1)$$

Algebraic Approach: Key Property

Key Property

If the message space W is chosen carefully, then the labelling φ can be made linear.

In general, W can be chosen as $R/(\pi_1) \times \cdots \times R/(\pi_k)$, where $\pi_1, \ldots, \pi_k$ are invariant factors of Λ/Λ'.

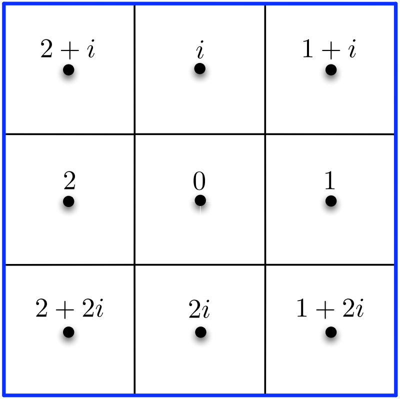

$$W = \mathbb{Z}[i]/(3)$$

$$\varphi(a+bi) = (a+bi) \bmod 3$$

Algebraic Approach: Encoding and Decoding

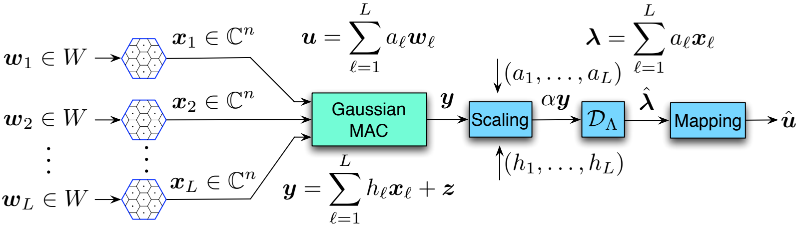

Encoding and Decoding

Transmitter ℓ sends $\mathbf{x}_\ell = \bar{\varphi}(\mathbf{w}_\ell) - Q_{\Lambda'}(\bar{\varphi}(\mathbf{w}_\ell))$

Receiver computes $\hat{\mathbf{u}} = \varphi(\mathcal{D}_\Lambda(\alpha \mathbf{y}))$

Remark: complexity here $\approx$ complexity for point-to-point channels

Algebraic Approach: Error Probability

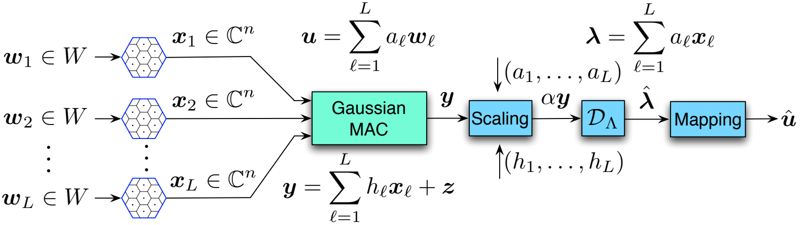

Error Probability

$$\Pr[\text{error}] = \Pr[\mathcal{D}_\Lambda(\mathbf{n}_{\text{eff}}) \notin \Lambda']$$

where $\mathbf{n}_{\text{eff}} \triangleq \sum_\ell (\alpha h_\ell - a_\ell)\mathbf{x}_\ell + \alpha \mathbf{z}$ is the effective noise.

Proof Sketch:

$\alpha \mathbf{y} = \sum_\ell a_\ell \mathbf{x}_\ell + \sum_\ell (\alpha h_\ell - a_\ell)\mathbf{x}_\ell + \alpha \mathbf{z} = \sum_\ell a_\ell \mathbf{x}_\ell + \mathbf{n}_{\text{eff}}$

Thus, $\mathcal{D}_\Lambda(\alpha \mathbf{y}) = \sum_\ell a_\ell \mathbf{x}_\ell + \mathcal{D}_\Lambda(\mathbf{n}_{\text{eff}})$

Therefore, $\hat{\mathbf{u}} = \varphi(\mathcal{D}_\Lambda(\alpha \mathbf{y})) = \mathbf{u} + \varphi(\mathcal{D}_\Lambda(\mathbf{n}_{\text{eff}}))$

Application: Practical Designs for Short Block Length

Union Bound Estimator

Union Bound Estimator (UBE) of the Error Probability

Recall that the effective noise $\mathbf{n}_{\text{eff}} = \sum_\ell (\alpha h_\ell - a_\ell)\mathbf{x}_\ell + \alpha \mathbf{z}$.

If Λ/Λ' admits hypercube shaping, then the UBE is

$$\Pr[\text{error}] \lesssim K(\Lambda/\Lambda') \exp\left(-\frac{d^2(\Lambda/\Lambda')}{4N_0(|\alpha|^2 + \text{SNR}\|\alpha\mathbf{h} - \mathbf{a}\|^2)}\right).$$

$K(\Lambda/\Lambda')$: # of the shortest vectors in the set difference $\Lambda - \Lambda'$

$d(\Lambda/\Lambda')$: length of the shortest vectors in the set difference $\Lambda - \Lambda'$

Implications

- minimize $K(\Lambda/\Lambda')$ and maximize $d(\Lambda/\Lambda')$
- minimize $Q(\alpha, \mathbf{a}) \triangleq |\alpha|^2 + \text{SNR}\|\alpha\mathbf{h} - \mathbf{a}\|^2$

Remark: $\mathbf{n}_{\text{eff}}$ has i.i.d. component with variance $N_0 Q(\alpha, \mathbf{a}) \Rightarrow$ minimum variance criterion

Figures of Merit

Signal-to-Effective-Noise Ratio (SENR)

$$\mathrm{SENR} \triangleq P/N_0 Q(\alpha, \mathbf{a})$$

Nominal Coding Gain

$$\Pr[\text{error}] \lesssim K(\Lambda/\Lambda') \exp\left(-\frac{3}{2}\frac{d^2(\Lambda/\Lambda')}{V(\Lambda)^{1/n}}\frac{\mathrm{SENR}}{2^{R_{\mathrm{mes}}}}\right).$$

Thus, $\gamma_c(\Lambda/\Lambda') \triangleq d^2(\Lambda/\Lambda')/V(\Lambda)^{1/n}$ is nominal coding gain

Effective Coding Gain

Rule of thumb: effective coding gain $=$ nominal coding gain
- 0.2dB $\times \log_2(K(\Lambda/\Lambda')/4)$

Practical Designs via Complex Construction A

Setup

- 9-QAM + linear codes over $\mathbb{Z}[i]/(3)$
- Idea: terminated (feedforward) convolutional codes
- Why? better performance-complexity tradeoff
- Low complexity $\Rightarrow$ constraint length $\nu = 1$ or 2
- Block length $= 200$

ν	$\mathbf{g}(D)$	$\gamma_c(\Lambda/\Lambda')$
1	$[1+(1+i)D,\ (1+i)+D]$	2 (3 dB)
2	$[1+D+(1+i)D^2,\ (1+i)+(1-i)D+D^2]$	3 (4.77 dB)

Remark: $\nu = 1 \Rightarrow 9$ states; $\nu = 2 \Rightarrow 81$ states

Practical Designs via Complex Construction A