List Decoding of Direct Sum Codes

Fernando Granha Jeronimo [1]
joint work with
Vedat Levi Alev, Dylan Quintana, Shashank Srivastava and Madhur Tulsiani

SODA 2020

[1] Supported by NSF travel award.

Outline

Outline

Outline

Outline

Outline

Context

Context

Binary codes are not that well understood compared to larger alphabet codes. We lack:

- algorithmic tools,
- explicit constructions, and
- impossibility results.

Context

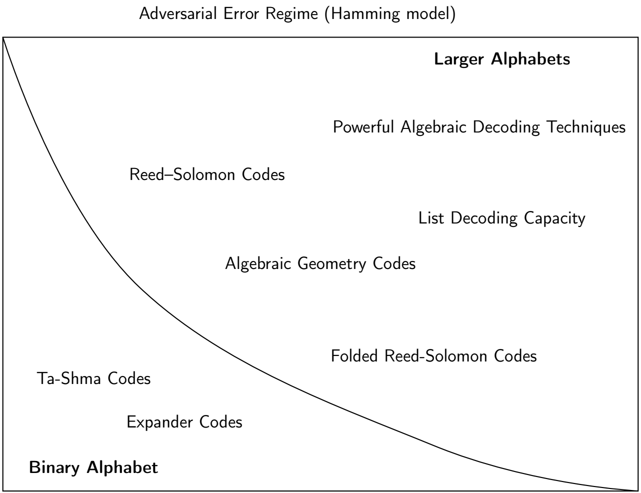
Adversarial Error Regime (Hamming model)
Larger Alphabets
Powerful Algebraic Decoding Techniques
Reed–Solomon Codes
List Decoding Capacity
Algebraic Geometry Codes
Folded Reed-Solomon Codes
Ta-Shma Codes
Expander Codes
Binary Alphabet

Context

Popular approach: obtain results by concatenating with binary codes

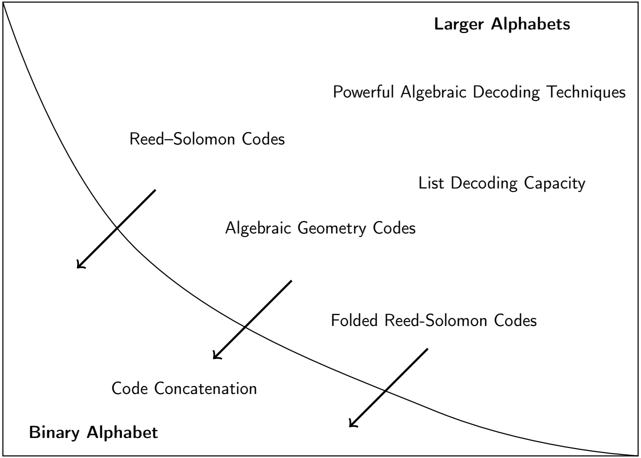

Context

This work

We make partial progress on the **algorithmic tool** front by providing a list decoding framework to handle *direct sum codes* on some (sparse) “expanding structures”.

Notation

Notation

- Let Σ be a finite alphabet.
- A code is a subset $\mathcal{C} \subseteq \Sigma^n$ (where n the block length).
- The (normalized) Hamming distance between $z, z' \in \Sigma^n$ is $\Delta(z, z') := |\{i \mid z_i \neq z'_i\}|/n$.
- The distance $\Delta(\mathcal{C})$ of $\mathcal{C}$ is $\min_{z,z' \in \mathcal{C}\colon z \neq z'} \Delta(z, z')$.
- The rate $r(\mathcal{C})$ of $\mathcal{C}$ is $\log_{|\Sigma|}(\mathcal{C})/n$.

Notation

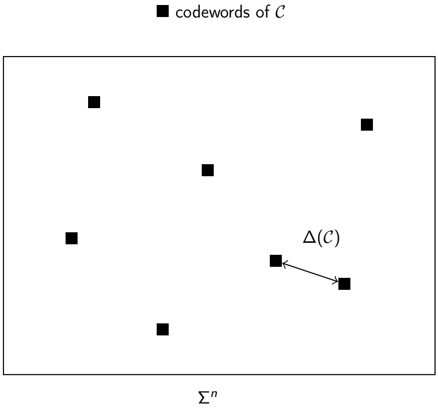

Notation

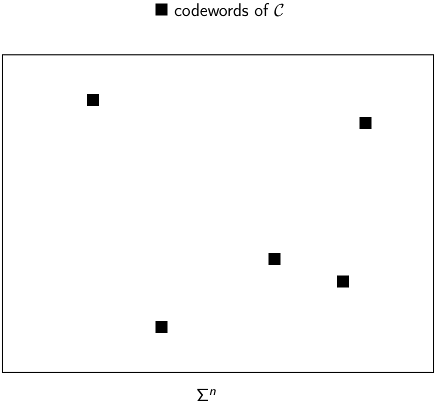

Lower rate $r(\mathcal{C})$

Notation

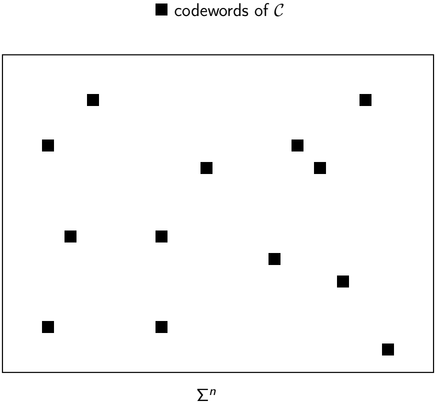

Higher rate $r(\mathcal{C})$

Background

We digress a bit to provide some background.

Expander and Codes

Expander graphs and codes have had a synergetic relationship. There are two major approaches:

- **Distance amplification:** use pseudorandom properties to boost distance ([ABNNR92], [AEL95], [GI03], [DHKNT19], etc).
- **Parity Check Matrix:** adjacency of a bipartite expander is used to define a parity check matrix ([Sipser–Spielman94], [Zémor01], LDPCs, etc).

Background

We digress a bit to provide some background.

Expander and Codes

Expander graphs and codes have had a synergetic relationship. There are two major approaches:

- **Distance amplification:** use pseudorandom properties to boost distance ([ABNNR92], [AEL95], [GI03], [DHKNT19], etc). **(this talk!)**
- **Parity Check Matrix:** adjacency of a bipartite expander is used to define a parity check matrix ([Sipser–Spielman94], [Zémor01], LDPCs, etc).

Expansion and Distance Amplification

Direct Product

Let $z \in \mathbb{F}_2^n$ and $X(k) \subseteq [n]^k$. The *direct product* of z is $y \in (\mathbb{F}_2^k)^{X(k)}$ defined as

$$y_{(i_1,\ldots,i_k)} = (z_{i_1}, \ldots, z_{i_k}),$$

for every $(i_1, \ldots, i_k) \in X(k)$.

Expansion and Distance Amplification

Direct Product

Let $z \in \mathbb{F}_2^n$ and $X(k) \subseteq [n]^k$. The *direct product* of z is $y \in (\mathbb{F}_2^k)^{X(k)}$ defined as

$$y_{(i_1,\ldots,i_k)} = (z_{i_1}, \ldots, z_{i_k}),$$

for every $(i_1, \ldots, i_k) \in X(k)$.

Shortcoming

Resulting alphabet is no longer binary.

Expansion and Distance Amplification

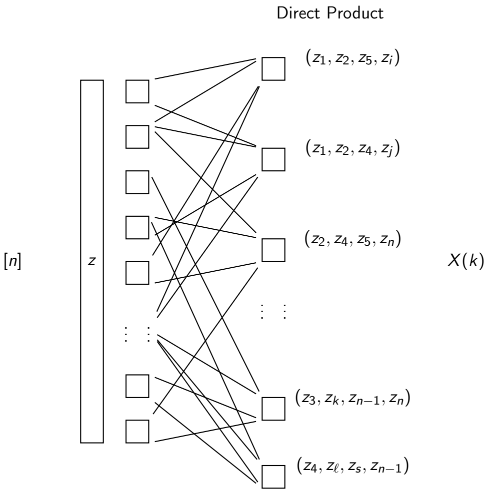

Expansion and Distance Amplification

Sampler [ABNNR92]

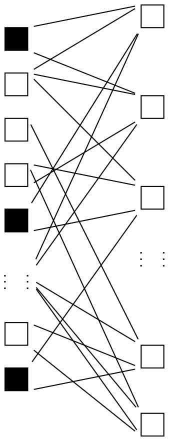

$[n]$ δ fraction of marked vertices on the left

Most vertices on the right have $\approx \delta$ fraction of marked neighbors $X(k)$

Expansion and Distance Amplification

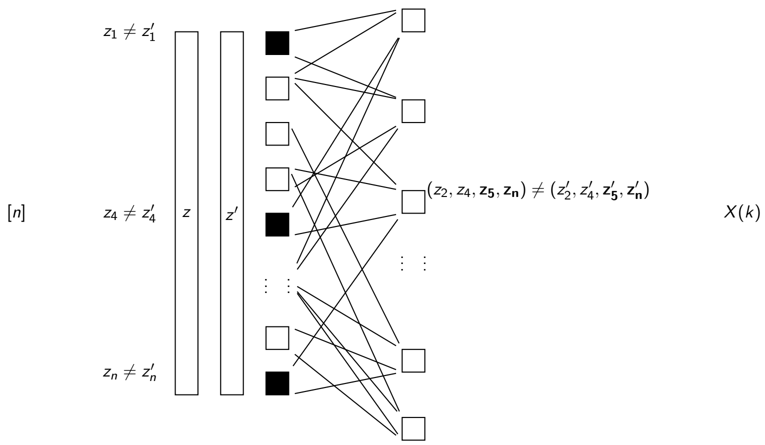

Expansion and Distance Amplification

Theorem (Dinur, Harsha, Kaufman, Navon and Ta-Shma'19)

For every $\beta > 0$, there is a family of explicit (non-binary) ***direct product*** *codes list decodable from radius $1 - \beta$ with rate $\exp(-\exp(\mathrm{poly}(1/\beta)))$ in polynomial time. (The construction relies on "**double samplers**").*

Expansion and Distance Amplification

Direct Sum

Let $z \in \mathbb{F}_2^n$ and $X(k) \subseteq [n]^k$. The *direct sum* of z is $y \in \mathbb{F}_2^{X(k)}$ defined as

$$y_{(i_1,\ldots,i_k)} = z_{i_1} \oplus \cdots \oplus z_{i_k},$$

for every $(i_1, \ldots, i_k) \in X(k)$. We denote $y = \mathrm{dsum}_{X(k)}(z)$.

Expansion and Distance Amplification

Direct Sum

Let $z \in \mathbb{F}_2^n$ and $X(k) \subseteq [n]^k$. The *direct sum* of z is $y \in \mathbb{F}_2^{X(k)}$ defined as

$$y_{(i_1,\dots,i_k)} = z_{i_1} \oplus \cdots \oplus z_{i_k},$$

for every $(i_1, \dots, i_k) \in X(k)$. We denote $y = \mathrm{dsum}_{X(k)}(z)$.

Advantage

The resulting alphabet is binary.

Expansion and Distance Amplification

Direct Sum

Let $z \in \mathbb{F}_2^n$ and $X(k) \subseteq [n]^k$. The *direct sum* of z is $y \in \mathbb{F}_{\mathbf{2}}^{\mathbf{X(k)}}$ defined as

$$y_{(i_1,\ldots,i_k)} = \mathbf{z_{i_1}} \oplus \cdots \oplus \mathbf{z_{i_k}},$$

for every $(i_1, \ldots, i_k) \in X(k)$. We denote $y = \mathrm{dsum}_{X(k)}(z)$.

Advantage

The resulting alphabet is binary.

Further Motivation

Ta-Shma [Ta-Shma'17] found explicit binary codes "near" the Gilbert–Varshamov bound using direct sum.

Expansion and Distance Amplification

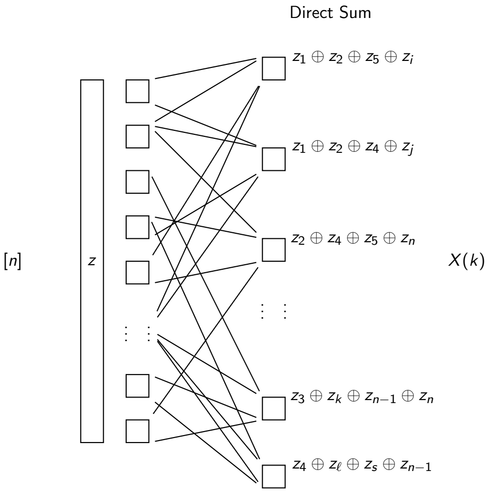

Expansion and Distance Amplification

Further Notation

- Let $z \in \mathbb{F}_2^n$. Define $\text{bias}(z) := |\mathbf{E}_{i \in [n]} (-1)^{z_i}|$.
- Let $\mathcal{C} \subseteq \mathbb{F}_2^n$. Define $\text{bias}(\mathcal{C}) := \max_{z \in \mathcal{C} \setminus \{0\}} \text{bias}(z)$.

Definition

Let $X \subseteq [n]^k$. We say that dsum_X is (β_0, β)-**parity sampler** iff

$$(\forall z \in \mathbb{F}_2^n) \left(\text{bias}(z) \leq \beta_0 \implies \text{bias}(\text{dsum}_X(z)) \leq \beta \right).$$

Expansion and Distance Amplification

A Dream Parity Sampler

Let $z \in \mathbb{F}_2^n$ with $\text{bias}(z) \leq \beta_0 < 1$. Let $X(k) = [n]^k$ (i.e., all k-tuples of $[n]$). Then

$$\text{bias}\left(\text{dsum}_{X(k)}(z)\right) \leq |\mathbf{E}_{i \in [n]}(-1)^{z_i}|^k \leq \beta_0^k.$$

Expansion and Distance Amplification

A Dream Parity Sampler

Let $z \in \mathbb{F}_2^n$ with $\text{bias}(z) \le \beta_0 < 1$. Let $X(k) = [n]^k$ (i.e., all k-tuples of $[n]$). Then

$$\text{bias}\left(\text{dsum}_{X(k)}(z)\right) \le \left|\mathbf{E}_{i \in [n]}(-1)^{z_i}\right|^k \le \beta_0^k.$$

Issue

$X(k)$ is **"too dense"**. The code $\text{dsum}_{X(k)}(\mathcal{C})$ has vanishing rate.

Expansion and Distance Amplification

Explicit Sparse Expanding Structures

Two explicit **sparse** "expanding" structures for parity sampling:

- (sparse) **High-dimensional expanders** [this work], and
- $X(k) \subseteq [n]^k$ from length-$(k-1)$ **walks on expander graph** $G = ([n], E)$ [Ta-Shma'17].

Expansion and Distance Amplification

High-dimensional expander (informal) definition

A γ-spectral high-dimensional expander X is a hypergraph s.t.

- **Downward-close**: if $\mathfrak{s} \in X$ and $\mathfrak{t} \subseteq \mathfrak{s}$, then $\mathfrak{t} \in X$.
- "**Multiscale expansion**": all sub-expander graphs are γ-two sided spectral expanders.

Expansion and Distance Amplification

Intuition

- Expander graph: "sparse approximation" of a complete graph
- High-dimensional expander: "sparse approximation" of $\binom{[n]}{\leq d}$

Expansion and Distance Amplification

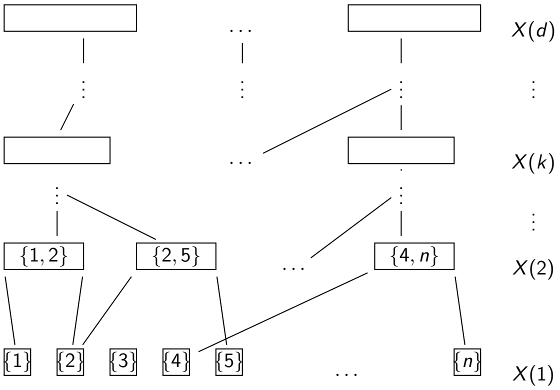

High-dimensional exapnder adjacency by containment.

Expansion and Distance Amplification

High-dimensional Expander as Parity Sampler

- High-dimensional expander $\approx$ complete hypergraph $\binom{[n]}{\leq d}$.
- Complete hypergraph $\binom{[n]}{\leq d}$ is a parity sampler.
- Follows that high-dimensional expander is a parity sampler.

Main Results

Theorem (Direct Sum High-dimensional Exapnders)

For every $\beta > 0$, there is a family of explicit ***binary*** *direct sum codes based on high-dimensional expanders list decodable from radius $1/2 - \beta$ with rate $\exp(-\mathrm{poly}(1/\beta))$ in time $n^{\mathrm{poly}(1/\beta)}$, where n is the block length.*

Main Results

Theorem (Direct Sum High-dimensional Exapnders)

For every $\beta > 0$, there is a family of explicit ***binary*** *direct sum codes based on high-dimensional expanders list decodable from radius $1/2 - \beta$ with rate $\exp(-\text{poly}(1/\beta))$ in time $n^{\text{poly}(1/\beta)}$, where n is the block length.*

Corollary (Direct Product High-dimensional Exapnders)

For every $\beta > 0$, there is a family of explicit (non-binary) direct product codes based on high-dimensional expanders list decodable from radius $1 - \beta$ with rate $\exp(-\text{poly}(1/\beta))$ in time $n^{\text{poly}(1/\beta)}$, where n is the block length.

Main Results

Theorem (Direct Sum Expander Walks)

*For every $\beta > 0$, there is a family of explicit **binary** direct sum codes based on walks on expanders list decodable from radius $1/2 - \beta$ with (**quasipolynomial**) rate exp$(-\mathrm{polylog}(1/\beta))$ in time $n^{\mathrm{poly}(1/\beta)}$, where n is the block length.*

Bird's-eye view of Techniques

Very High-level Strategy

- Start with a unique decoding algorithm for direct sum codes.
- Enhance this algorithm with list decoding capabilities.

Bird's-eye view of Techniques

Very High-level Strategy

- Start with a unique decoding algorithm for direct sum codes.
- Enhance this algorithm with list decoding capabilities.

First Step: Dealing with k-XOR

We first describe this unique decoding algorithm and how it is naturally related to k-XOR.

Bird's-eye view of Techniques: Unique Decoding

Setup

- $\mathcal{C} \subseteq \mathbb{F}_2^n$ a β_0-biased code,
- $X \subseteq [n]^k$ for direct sum, and
- $\mathcal{C}' = \text{dsum}_X(\mathcal{C})$ a β-biased code.

Bird's-eye view of Techniques: Unique Decoding

Suppose $y^* \in \mathcal{C}'$ is corrupted into some $\tilde{y} \in \mathbb{F}_2^X$ in the unique decoding ball centered at y^*.

Unique Decoding Scenario: k-XOR

Unique decoding $\tilde{y}$ amounts to solving

$$\arg\max_{z \in \mathcal{C}} \mathbf{E}_{(i_1,\ldots,i_k)\in X} \mathbf{1}[z_{i_1} \oplus \cdots \oplus z_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}],$$

which is a MAX k-XOR instance $\mathfrak{I}$ with the additional constraint that the solution z must lie in $\mathcal{C}$.

Bird's-eye view of Techniques: Unique Decoding

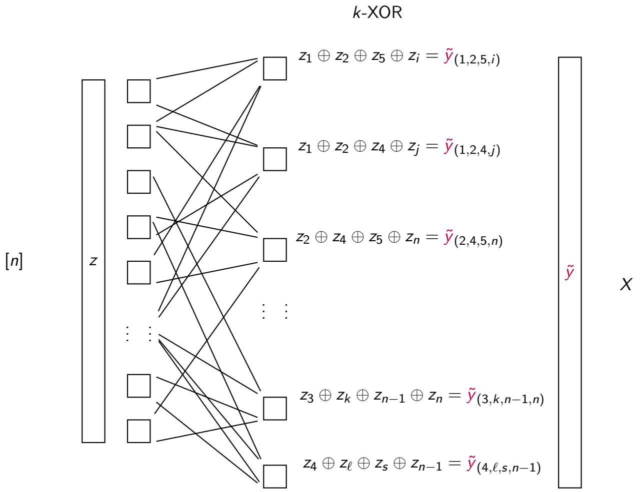

Bird's-eye view of Techniques: Unique Decoding

Let $z^* \in \mathcal{C}$ be s.t. $y^* = \mathrm{dsum}_X(z^*)$.

Optimal Value

Since $\tilde{y}$ is in the unique decoding ball centered at y^*, we have

$$\mathbf{E}_{(i_1,\dots,i_k)\in X}\mathbf{1}[{z^*}_{i_1} \oplus \cdots \oplus {z^*}_{i_k} \neq \tilde{y}_{(i_1,\dots,i_k)}] = \Delta(y^*, \tilde{y}) < \Delta(\mathcal{C}')/2$$

Thus,

$$\mathrm{OPT}(\mathfrak{I}) \geq \mathbf{E}_{(i_1,\dots,i_k)\in X}\mathbf{1}[{z^*}_{i_1} \oplus \cdots \oplus {z^*}_{i_k} = \tilde{y}_{(i_1,\dots,i_k)}] > 1 - \Delta(\mathcal{C}')/2$$

Bird's-eye view of Techniques: Unique Decoding

Optimal Solution

Suppose that we can find $\tilde{z} \in \mathbb{F}_2^n$ (rather than in $\mathcal{C}$) satisfying

$$\mathbf{E}_{(i_1,\ldots,i_k)\in X}\mathbf{1}[\tilde{z}_{i_1} \oplus \cdots \oplus \tilde{z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}] = \mathsf{OPT}(\mathfrak{I}) > 1 - \Delta(\mathcal{C}')/2$$

Thus, $\Delta(\mathrm{dsum}_X(\tilde{z}), \tilde{y}) < \Delta(\mathcal{C}')/2$

Bird's-eye view of Techniques: Unique Decoding

By triangle inequality,

$$\begin{aligned}\Delta(\mathrm{dsum}_X(\tilde{z}), \mathrm{dsum}_X(z^*)) \ \leq \ & \Delta(\mathrm{dsum}_X(\tilde{z}), \tilde{y}) + \\ & \Delta(\tilde{y}, \mathrm{dsum}_X(z^*)) < \Delta(\mathcal{C}') \leq 1/2 - \beta/2,\end{aligned}$$

implying

$$\mathrm{bias}(\mathrm{dsum}_X(\tilde{z}) \oplus \mathrm{dsum}_X(z^*)) = \mathrm{bias}(\mathrm{dsum}_X(\tilde{z} \oplus z^*)) > \beta$$

"Nontrivial bias"

Bird's-eye view of Techniques: Unique Decoding

Claim

If dsum_X is a "strong enough" parity sampler, then either $\tilde{z}$ or $\tilde{z} \oplus 1$ lie in the unique decoding ball of $\mathcal{C}$ centered at z^*.

Bird's-eye view of Techniques: Unique Decoding

Claim

If dsum_X is a $(1/2 + \beta_0/2, \beta)$-parity sampler, then either $\tilde{z}$ or $\tilde{z} \oplus 1$ lie in the unique decoding ball of $\mathcal{C}$ centered at z^*.

Proof

Towards a contradiction, suppose

$$\Delta(\mathcal{C})/2 \leq \Delta(\tilde{z}, z^*) \leq 1 - \Delta(\mathcal{C})/2,$$

implying that $\mathrm{bias}(\tilde{z} \oplus z^*) \leq 1 - \Delta(\mathcal{C}) \leq 1/2 + \beta_0/2$. "not too large"
Using the $(1/2 + \beta_0/2, \beta)$-parity sampler assumption,

$$\mathrm{bias}(\mathrm{dsum}_X(\tilde{z} \oplus z^*)) \leq \beta, \qquad \text{"small"}$$

contradicting $\mathrm{bias}(\mathrm{dsum}_X(\tilde{z} \oplus z^*)) > \beta$ "Nontrivial bias" from before.

Bird's-eye view of Techniques: Unique Decoding

Moral

- Find solution $\tilde{z} \in \mathbb{F}_2^n$ (rather than in $\mathcal{C}$) is enough.
- Unique decoder of $\mathcal{C}$: correct $\tilde{z}$ into z^*.

Bird's-eye view of Techniques: Unique Decoding

Need to resolve the following assumption.

Optimal Solution

Suppose that we can find $\tilde{z} \in \mathbb{F}_2^n$ (rather than $\tilde{z} \in \mathcal{C}$) satisfying

$$\mathbf{E}_{(i_1,\ldots,i_k)\in X} \mathbf{1}[\tilde{z}_{i_1} \oplus \cdots \oplus \tilde{z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}] = \mathsf{OPT}(\mathfrak{I})$$

Bird's-eye view of Techniques: Unique Decoding

Need to resolve the following assumption.

Optimal Solution

Suppose that we can find $\tilde{z} \in \mathbb{F}_2^n$ (rather than $\tilde{z} \in \mathcal{C}$) satisfying

$$\mathbf{E}_{(i_1,\ldots,i_k)\in X}\mathbf{1}[\tilde{z}_{i_1} \oplus \cdots \oplus \tilde{z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}] = \mathsf{OPT}(\mathfrak{I})$$

Possible issue?

MAX k-XOR is NP-hard, right?

Bird's-eye view of Techniques: Unique Decoding

Possible issue?

MAX k-XOR is NP-hard, right?

Not an issue

Right, it can be NP-hard in general. However, for **expanding** instances we can find an **approximate** solution (and that is enough).

Bird's-eye view of Techniques: Unique Decoding

Not an issue

Right, it can be NP-hard in general. However, for **expanding** instances we can find an **approximate** solution (and that is enough).

Using a better spectral analysis of Dinur–Dikstein'19.

Theorem (Alev–J–Tulsiani'19)

Let X be a γ-spectral high-dimensional expander on n vertices. Let $\mathfrak{I}$ be a k-CSP on $X(k)$ with alphabet size q.
If $\gamma \leq \mathrm{poly}(\epsilon/q^k)$, then we can find a solution $z \in [q]^n$ satisfying

$$\mathrm{OPT}(\mathfrak{I}) - \epsilon,$$

fraction of the constraints of $\mathfrak{I}$ in time $n^{\mathrm{poly}(q^k/\epsilon)}$.

Bird's-eye view of Techniques: Unique Decoding

Theorem (this work)

Alev–J–Tulsiani'19 also holds when $X(k)$ is the collection of length-$(k-1)$ walks on a γ-two-sided spectral expander graph G.

Bird's-eye view of Techniques: List Decoding

List Decoding

How about our main goal of list decoding?

Bird's-eye view of Techniques: List Decoding

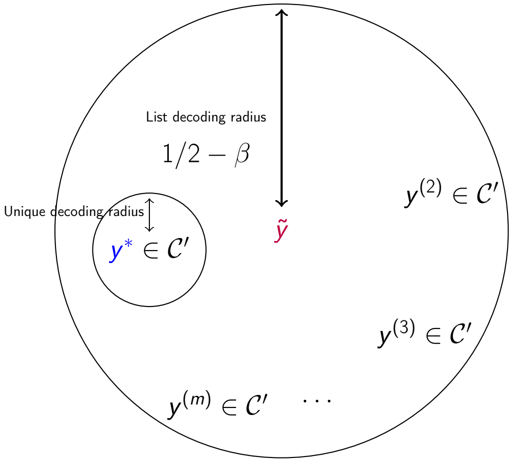

Bird's-eye view of Techniques: List Decoding

List Decoding Task

Given $\tilde{y}$ promised to satisfy $\Delta(\tilde{y}, \mathcal{C}') \leq 1/2 - \beta$, we want to find

$$\mathcal{L}(\tilde{y}) \coloneqq \left\{ y \in \mathcal{C}' \mid \Delta(y, \tilde{y}) \leq \frac{1}{2} - \beta \right\}.$$

Bird's-eye view of Techniques: List Decoding

CSP Algorithms

We will need to understand a bit more the preceding CSP algorithms for expanding structures which are based on the *Sum-of-Squares* (SOS) semidefinite programming hierarchy.

Bird's-eye view of Techniques: List Decoding

Issue: "Low Entropy" convex program solution

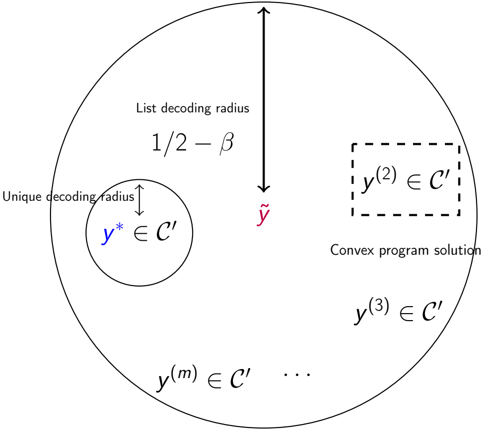

Bird's-eye view of Techniques: List Decoding

Want: "High Entropy" convex program solution

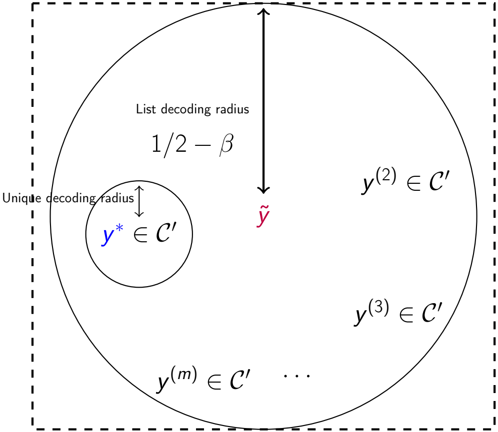

Convex program solution

Bird's-eye view of Techniques: List Decoding

Sum-of-Square Solution: PSD ensemble

A t-local PSD ensemble ensemble is a collection $\mathbf{Z}_1, \ldots, \mathbf{Z}_n$ of "local random variables" taking value in $\{\pm 1\}$ and satisfying:

- for $S \subseteq [n]$ with $|S| \leq t$, the variable $\mathbf{Z}_S$ has a distribution μ_S.
- for $S, T \subseteq [n]$ with $|S|, |T| \leq t$, $\mu_{S|T} = \mu_{T|S}$.
- a global PSD property (we won't have time to describe).

Bird's-eye view of Techniques: List Decoding

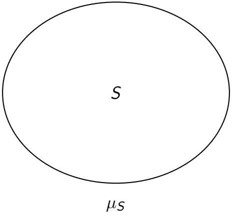

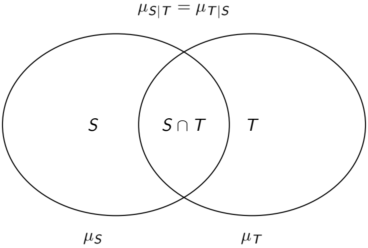

Bird's-eye view of Techniques: List Decoding

Pseudo-expectation

$\widetilde{\mathbf{E}}$ will denote the "expectation" w.r.t. the local random variables.

Bird's-eye view of Techniques: List Decoding

$$\begin{aligned} &\text{maximize} \quad \mathbf{E}_{(i_1,\ldots,i_k)\in X}\widetilde{\mathbf{E}}\left[\mathbf{1}[\mathbf{Z}_{i_1}\cdots\mathbf{Z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}]\right] \quad \text{(Objective)} \\ &\text{subject to} \\ &\quad \mathbf{Z}_1,\ldots,\mathbf{Z}_n \text{ being } t\text{-local PSD ensemble} \end{aligned}$$

Table: k-XOR unique decoding SOS formulation for $\tilde{y}$.

Bird's-eye view of Techniques: List Decoding

Recall the SOS program. Do you see the issue for list decoding?

$$\begin{aligned} &\text{maximize} && \mathbf{E}_{(i_1,\ldots,i_k)\in X}\widetilde{\mathbf{E}}\left[\mathbf{1}[\mathbf{Z}_{i_1}\cdots\mathbf{Z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}]\right] \quad \text{(Objective)} \\ &\text{subject to} && \\ & && \mathbf{Z}_1,\ldots,\mathbf{Z}_n \text{ being } t\text{-local PSD ensemble} \end{aligned}$$

Table: k-XOR unique decoding SOS formulation for $\tilde{y}$.

Bird's-eye view of Techniques: List Decoding

Recall the SOS program. Do you see the issue for list decoding?

$$\begin{array}{ll} \text{maximize} & \mathbf{E}_{(i_1,\ldots,i_k)\in X}\widetilde{\mathbf{E}}\left[\mathbf{1}[\mathbf{Z}_{i_1}\cdots\mathbf{Z}_{i_k} = \tilde{y}_{(i_1,\ldots,i_k)}]\right] \quad \text{(Objective)} \\ \text{subject to} & \\ & \mathbf{Z}_1,\ldots,\mathbf{Z}_n \text{ being } t\text{-local PSD ensemble} \end{array}$$

Table: k-XOR unique decoding SOS formulation for $\tilde{y}$.

Issue

The objective function "forces" the solution to agree with $\tilde{y}$ as much as possible.

Bird's-eye view of Techniques: List Decoding

Direct Sum of the PSD ensemble

Let $X \subseteq [n]^k$. For $\mathfrak{s} = (i_1, \ldots, i_k) \in X$, define the local random variable (we are working with $\{\pm 1\}$ variables now)

$$\mathbf{Y}_{\mathfrak{s}} := \mathbf{Z}_{i_1} \cdots \mathbf{Z}_{i_k}.$$

$\{Y_{\mathfrak{s}}\}_{\mathfrak{s} \in X}$ is also a PSD ensemble.

Bird's-eye view of Techniques: List Decoding

List Decoding Attempt

Drop the objective function and add a constraint?

$$\mathbf{E}_{\mathfrak{s}\in X(k)}\widetilde{\mathbf{E}}\left[\tilde{y}_{\mathfrak{s}} \cdot \mathbf{Y}_{\mathfrak{s}}\right] \geq 2\beta \qquad \text{(Agreement Constraint)}$$

$$\mathbf{Y}_{\mathfrak{s}} := \mathbf{Z}_{i_1} \cdots \mathbf{Z}_{i_k} \qquad (\forall \mathfrak{s} = (i_1, \ldots, i_k) \in X(k))$$

$\mathbf{Z}_1, \ldots, \mathbf{Z}_n$ being t-local PSD ensemble

Bird's-eye view of Techniques: List Decoding

List Decoding Attempt

Drop the objective function and add a constraint?

$$\mathbf{E}_{\mathfrak{s} \in X(k)} \widetilde{\mathbf{E}} \left[\tilde{y}_{\mathfrak{s}} \cdot \mathbf{Y}_{\mathfrak{s}} \right] \geq 2\beta \qquad \text{(Agreement Constraint)}$$

$$\mathbf{Y}_{\mathfrak{s}} := \mathbf{Z}_{i_1} \cdots \mathbf{Z}_{i_k} \qquad (\forall \mathfrak{s} = (i_1, \dots, i_k) \in X(k))$$

$\mathbf{Z}_1, \dots, \mathbf{Z}_n$ being t-local PSD ensemble

Issue

SOS solution may not be "diverse" enough. In particular, a delta distribution with any single element in $\mathcal{L}(\tilde{y})$ is a feasible solution.

Bird's-eye view of Techniques: List Decoding

Issue

SOS solution may not be "diverse" enough. In particular, a delta distribution with any single element in $\mathcal{L}(\tilde{y})$ is a feasible solution.

The goal

Make the SOS solution richer ("high entropy") somehow.

Bird's-eye view of Techniques: List Decoding

Solution: use a proxy for negative entropy to enforce diversity in the SOS solution.

Definition (Entropic Proxy)

Let $\mathbf{Y} = \{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s} \in X(k)}$ be a t-local PSD ensemble. Define $\Psi = \Psi\left(\{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s} \in X(k)}\right)$ as

$$\Psi \ := \ \mathbf{E}_{\mathfrak{s},\mathfrak{t} \in X(k)} \left(\widetilde{\mathbf{E}}\left[\mathbf{Y}_{\mathfrak{s}} \mathbf{Y}_{\mathfrak{t}}\right]\right)^2 .$$

Bird's-eye view of Techniques: List Decoding

Solution: use a proxy for negative entropy to enforce diversity in the SOS solution.

Definition (Entropic Proxy)

Let $\mathbf{Y} = \{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s} \in X(k)}$ be a t-local PSD ensemble. Define $\Psi = \Psi\left(\{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s} \in X(k)}\right)$ as

$$\Psi := \mathbf{E}_{\mathfrak{s},\mathfrak{t} \in X(k)} \left(\widetilde{\mathbf{E}}\left[\mathbf{Y}_{\mathfrak{s}} \mathbf{Y}_{\mathfrak{t}}\right]\right)^2.$$

A similar idea was independently used by Raghavendra–Yau'19 and Karmalkar–Klivans–Kothari'19 both in the context of learning regression.

Bird's-eye view of Techniques: List Decoding

minimize	$\Psi\left(\{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s}\in X(k)}\right)$	(Negative Entropy Proxy)
subject to		
	$\mathbb{E}_{\mathfrak{s}\in X(k)}\widetilde{\mathbb{E}}\left[\tilde{y}_{\mathfrak{s}}\cdot\mathbf{Y}_{\mathfrak{s}}\right]\geq 2\beta$	(Agreement Constraint)
	$\mathbf{Y}_{\mathfrak{s}} := \mathbf{Z}_{i_1}\cdots\mathbf{Z}_{i_k}$	$(\forall\mathfrak{s}=(i_1,\ldots,i_k)\in X(k))$
	$\mathbf{Z}_1,\ldots,\mathbf{Z}_n$ being t-local PSD ensemble	

Table: List decoding SOS formulation for $\tilde{y}$.

Bird's-eye view of Techniques: List Decoding

Why does Ψ enforce diversity in the SOS solution?

If SOS solution contains a single codeword

$$\begin{array}{|c|c|c|c|} \hline 1 & -1 & 1 & \cdots \quad y^{(i)} \qquad\qquad\qquad \\ \hline \end{array}$$

$\mathfrak{s}$ $\mathfrak{t}$

$y^{(i)} \in \mathcal{L}(\tilde{y})$

$$\Psi = \mathbf{E}_{\mathfrak{s},\mathfrak{t} \in X(k)} \left(\widetilde{\mathbf{E}} \left[\mathbf{Y}_{\mathfrak{s}} \mathbf{Y}_{\mathfrak{t}} \right] \right)^2 = 1 \text{ (as large as possible)}$$

Bird's-eye view of Techniques: List Decoding

Why does Ψ enforce diversity in the SOS solution?

If SOS solution is uniform on two codewords

1	−1	1	⋯ $y^{(i)}$
1	1	−1	⋯ $y^{(j)}$
$\mathfrak{s}$	$\mathfrak{t}$		

$y^{(i)}, y^{(j)} \in \mathcal{L}(\tilde{y})$

$$\Psi = \mathbf{E}_{\mathfrak{s},\mathfrak{t} \in X(k)} \left(\widetilde{\mathbf{E}} \left[\mathbf{Y}_{\mathfrak{s}} \mathbf{Y}_{\mathfrak{t}} \right] \right)^2 < 1$$

Bird's-eye view of Techniques: List Decoding

Propagation Rounding Algorithm

- Choose $\ell \in [t/k]$.
- Sample $S \sim \binom{X}{\ell}$.
- Sample an assignment $\eta \sim \mathbf{Y}_S$.
- Sample $z_i \sim \{\mathbf{Z}_i | \mathbf{Y}_S = \eta\}$ independently for $i \in [n]$.
- Return assignment $(z_1, \ldots, z_n)$.

Bird's-eye view of Techniques: List Decoding

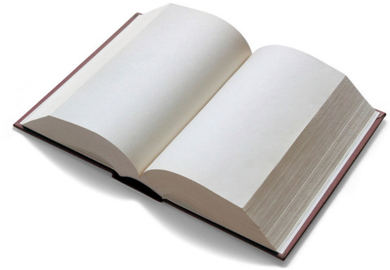

Figure: SOS solution is like a book. Each choice of S and η lead to a "page" (or set of pages), i.e., solution(s).

Bird's-eye view of Techniques: List Decoding

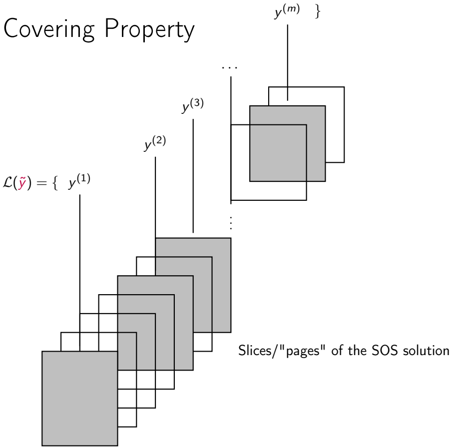

Questions

How far can we push this technique?

- Can we get rate $\Omega(\beta^{O(1)})$?
- Can we decode Ta-Shma's construction?
- Can we do better than the (algebraic) state-of-the-art rate $\Omega(\beta^3)$?

That's all.

Thank you!