Time and Space Efficient Deterministic Decoders

Joshua Cook, Dana Moshkovitz
The University of Texas at Austin

What is an Error Correcting code?

Definition (Error Correcting Code)**:**

Function C: $\Sigma^k \rightarrow \Sigma^n$ such that for $x \neq y$, we have

$$\Pr_i[C(x)_i \neq C(y)_i] \geq \boldsymbol{\delta}^*.$$

Relative distance $\boldsymbol{\delta}^*$, rate $r = k/n$. Good if $\boldsymbol{\delta}^*$, $r = \boldsymbol{\Omega}(1)$ and $|\Sigma| = O(1)$.

The correcting radius is $\boldsymbol{\delta} = \boldsymbol{\delta}^*/2$.

Many, many applications,

some of which are in the small space regime.

What is decoding?

Given w close to C(x), output x.

If w is within **δ** of C(x), we decode x.

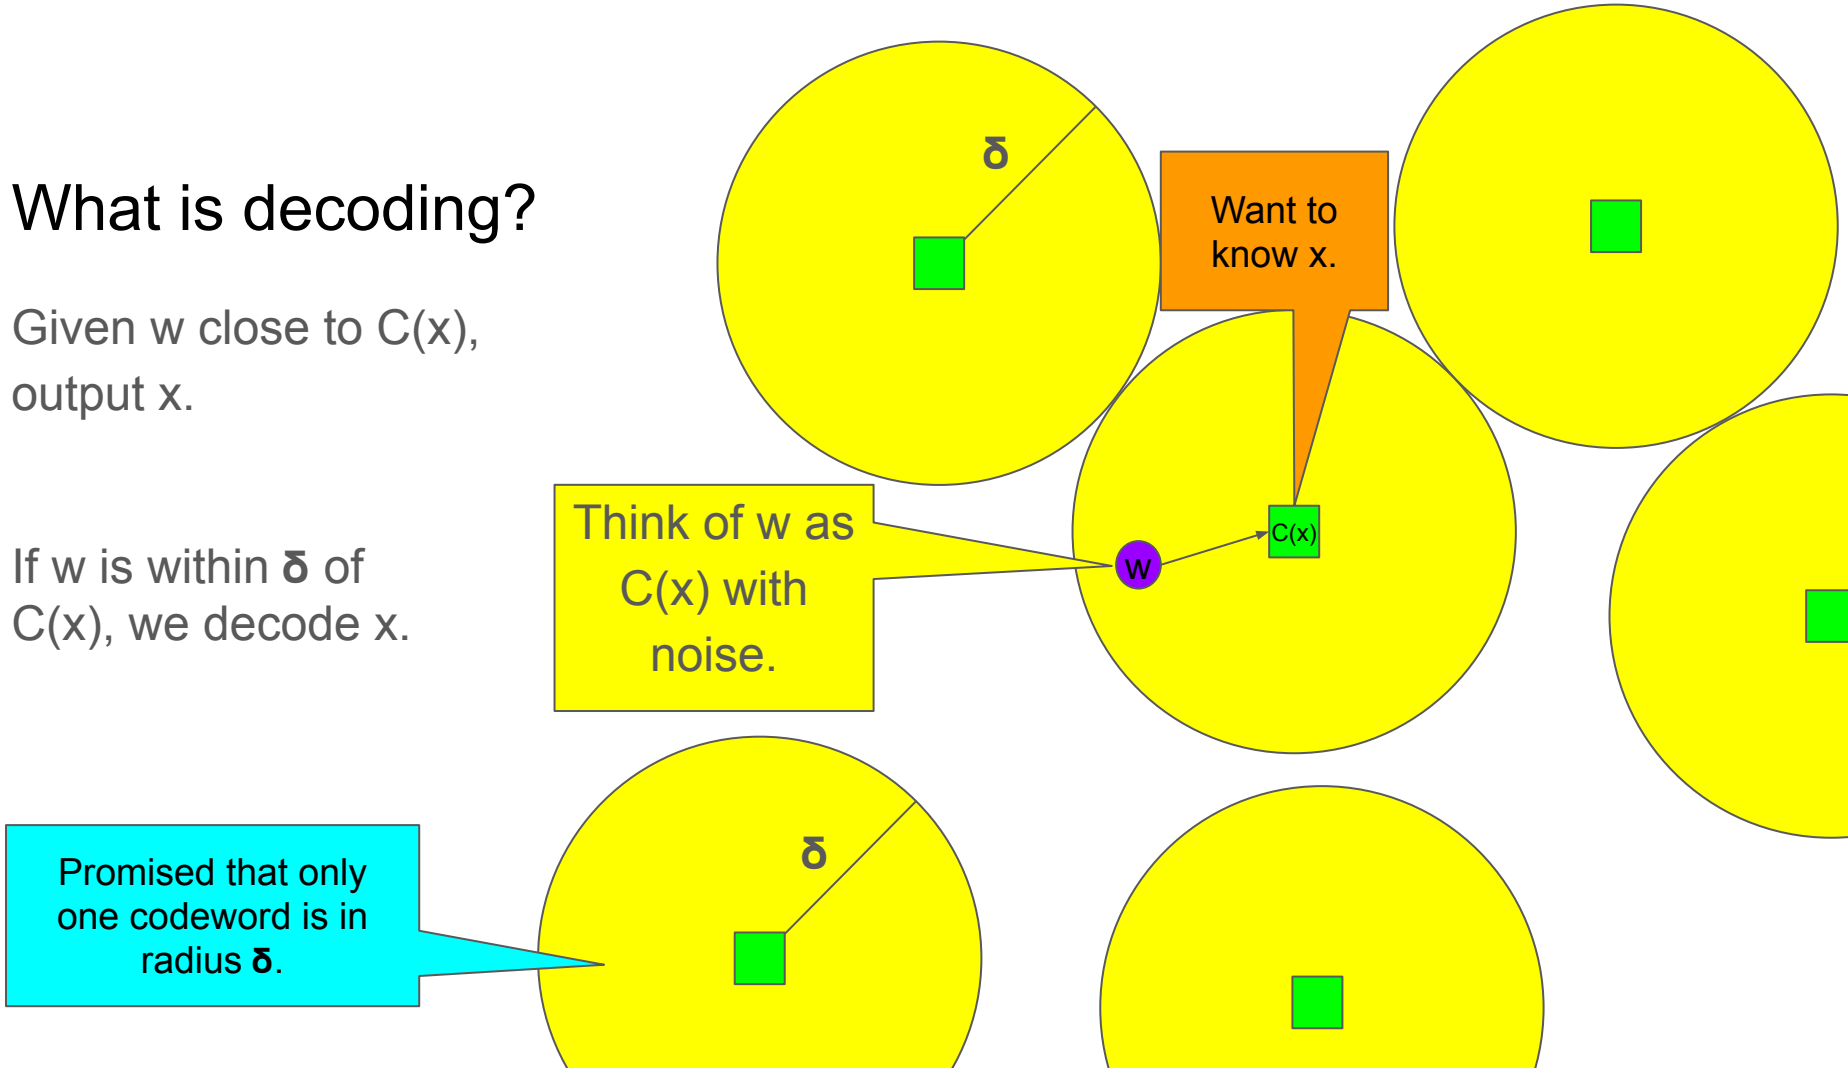

What is a low space decoder?

Input is read only. Output is write only. Have only $n^{o(1)}$ working space.

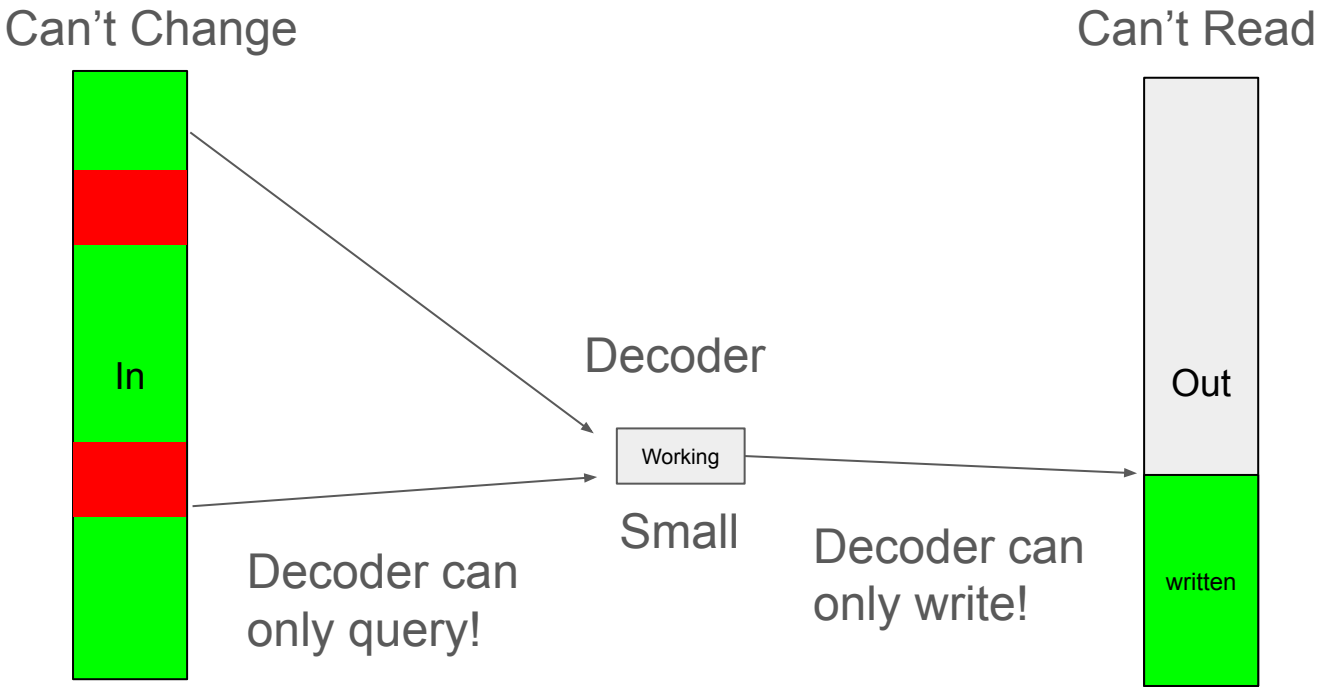

Reed-Solomon Codes

Reed-Solomon codes: univariate low degree polynomials.

Example: message $a = (a_1,a_2,a_3)$, field of order 5.

Let $f_a(x) = a_1 + a_2x + a_3x^2$. Codeword is: $[f_a(0), f_a(1), f_a(2), f_a(3), f_a(4)]$.

Has known (almost) linear time (and space) decoders.

Examples: [0, 1, 2, 3, 4] is the codeword of (0, 1, 0), $f(x) = x$.

[1, 2, 0, 0, 2] is the codeword of (1, 0, 1), $f(x) = 1 + x^2$.

Reed-Muller Codes

Reed-Muller codes: codewords are multivariate low degree polynomials.

$f(x,y) = a_1 + a_2x + a_3x^2 + a_4xy + a_5y + a_6y^2$.

f(0,0)	f(1,0)	f(2,0)	f(3,0)	
f(0,1)	f(1,1)	f(2,1)		f(4,1)
f(0,2)	f(1,2)		f(3,2)	f(4,2)
f(0,3)		f(2,3)	f(3,3)	f(4,3)
	f(1,4)	f(2,4)	f(3,4)	f(4,4)

g(0) g(1) g(2) g(3) g(4)

f restricted to line is univariate.
Example: g(x) = f(x,4-x)

Local Decoding

Correct a symbol with a small number of queries.

Take a random line through a point.

Works when line hits few corruptions.

Deterministic local correction?

No.

Always checks the same few symbols, adversary corrupts those.

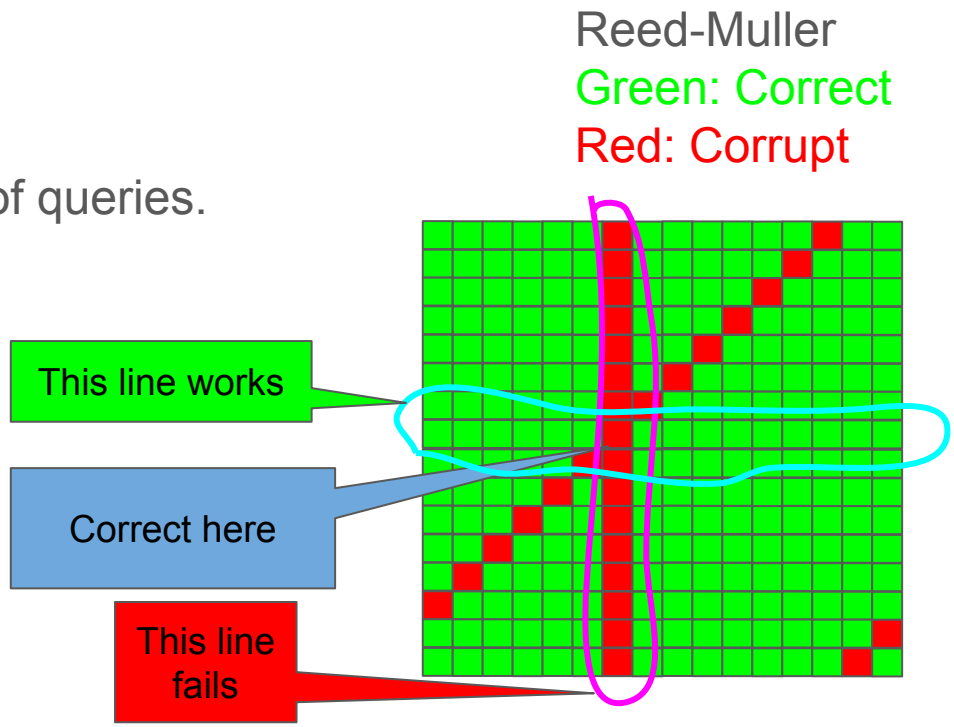

Locally Correctable Codes (LCC)

Definition (Locally Correctable Code (LCC))**:**

An LCC is a code $C: \Sigma^k \to \Sigma^n$ with a randomized algorithm D such that:

For any $w \in \Sigma^n$, $x \in \Sigma^k$ where $\Pr_i[C(x)_i \neq w_i] \leq \boldsymbol{\delta}'$ and any $i \in [n]$ we have

$$\Pr[D(w, i) \neq C(x)_i] \leq \tfrac{1}{3}$$

D is q query if it only makes q queries to w. $\boldsymbol{\delta}'$ is called the correcting radius.

Most codes we consider are systematic (the message is contained as plain text within the codeword), so local correctors are decoders.

Randomized Decoder Continued

Reed-Muller codes are locally correctable!

By repeating O(log(n)) times, the error probability drops below 1/n.

So it is unlikely any symbol is decoded incorrectly.

This gives a time and space efficient **randomized** decoder.

Local correction cannot be **deterministic**!

This decoder is **non-adaptive** (queries do not depend on the input).

Non-Adaptive, Deterministic Decoders Fail

Gronemeier: non-adaptive deterministic decoders can't be sub-polynomial space and almost linear time.

Non-adaptive means read and write locations are independent of the input.

Idea: for space S, wait for an interval where the decoder outputs S+1 symbols and only reads o(n) input symbols.

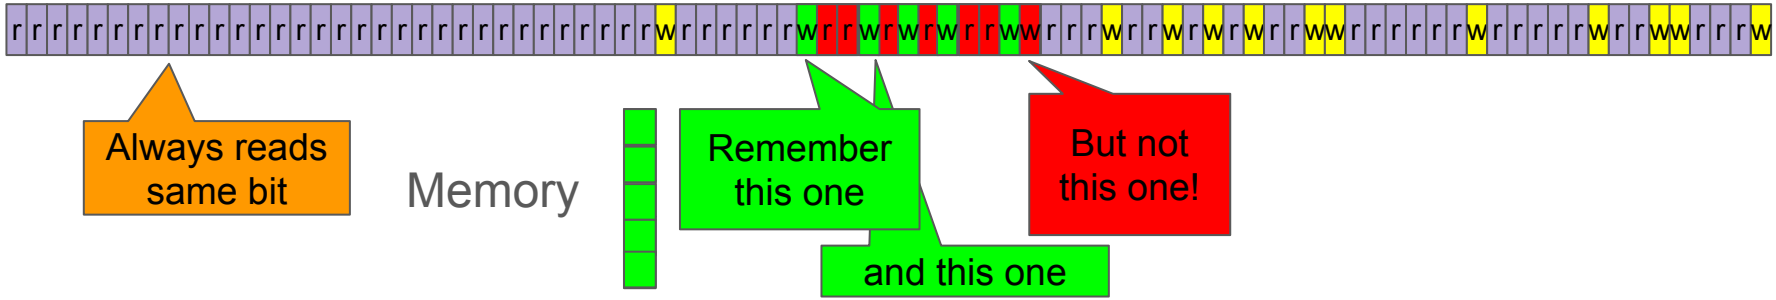

First Result: Efficient (Non-Uniform) Deterministic Decoders

Theorem (Deterministic Decoder for LCC)**:**

Good, typical* $q = n^{\alpha}$ query LCCs have **non-uniform**, *deterministic* decoders running in space $O(q \log(n)^2)$ and time $n^{1+O(\sqrt{\alpha})}$.

*(typical means systematic, non-adaptive, and with perfect completeness).

For $q = n^{o(1)}$: space is $n^{o(1)}$ and time $n^{1+o(1)}$.

Gronemeier proved non-adaptive decoders for good codes required time T and space S such that $ST = \Omega(n^2)$.

A flawed approach

Find a (single) q query f such that

$$\Pr_{i,j}[C(x)_i \neq f_j(w)_i] \leq \eta \Pr_i[C(x)_i \neq w_i]$$

f reduces corruptions by an η fraction.

Tradeoff between q and η.

Need both q and η small.

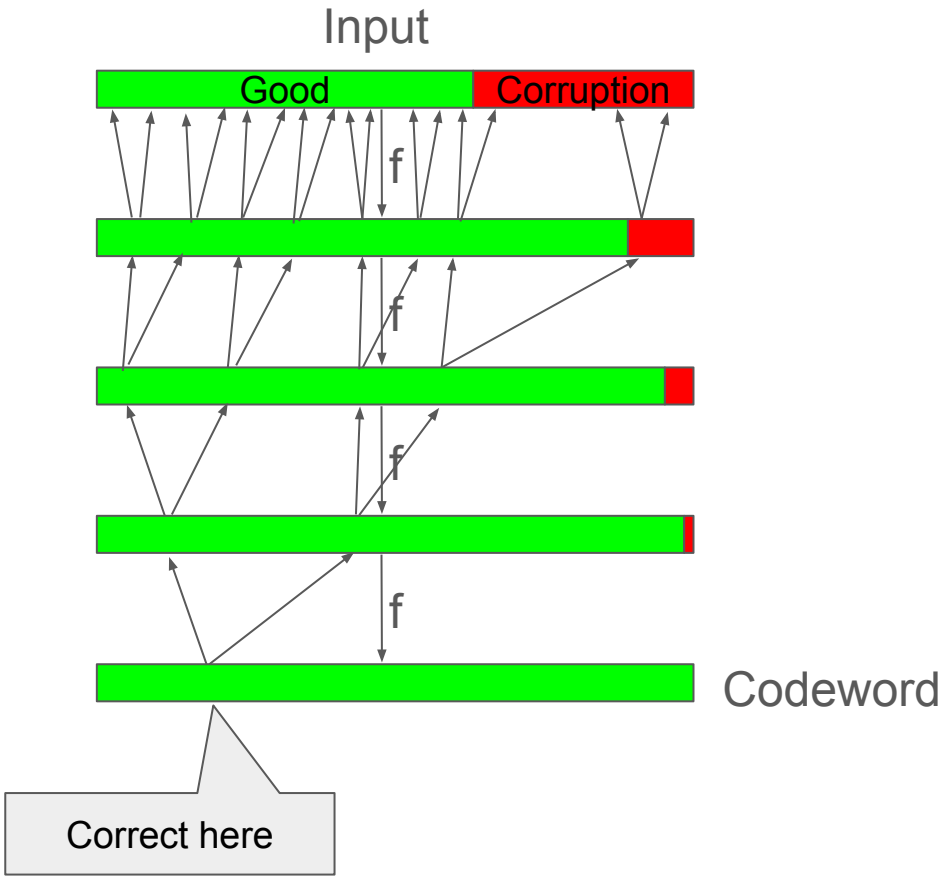

How good can a single, deterministic f be?

If f makes q queries, and we can corrupt **δ** fraction of symbols.

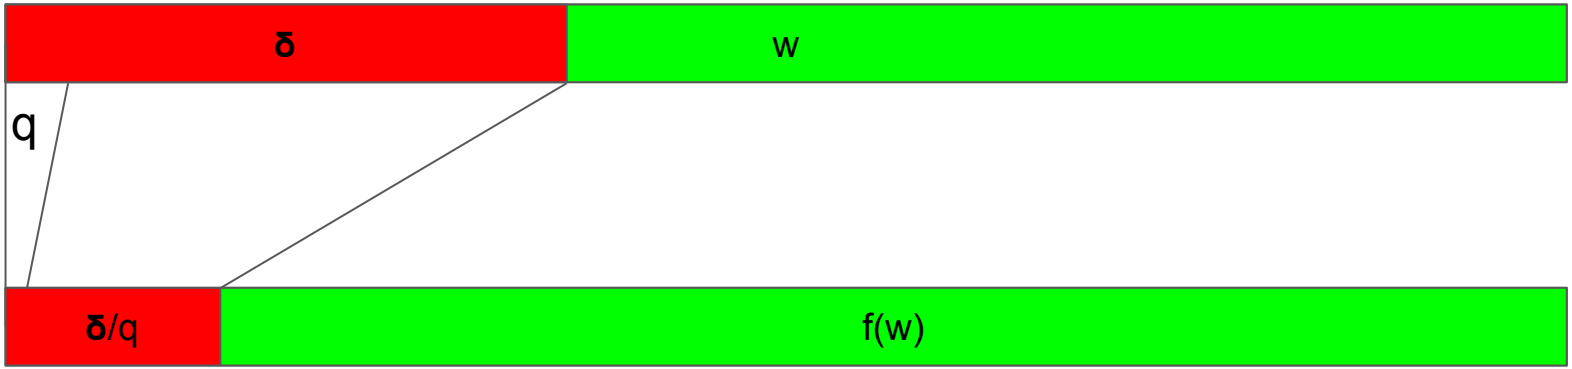

For deterministic f, only get η = 1/q

Why doesn’t it work?

Deterministic, q query f only reduces **δ** corruptions to **δ**/q.

To get zero errors, we would need **δ**n queries (per symbol)!

This gives a total time of **δ**n^2.

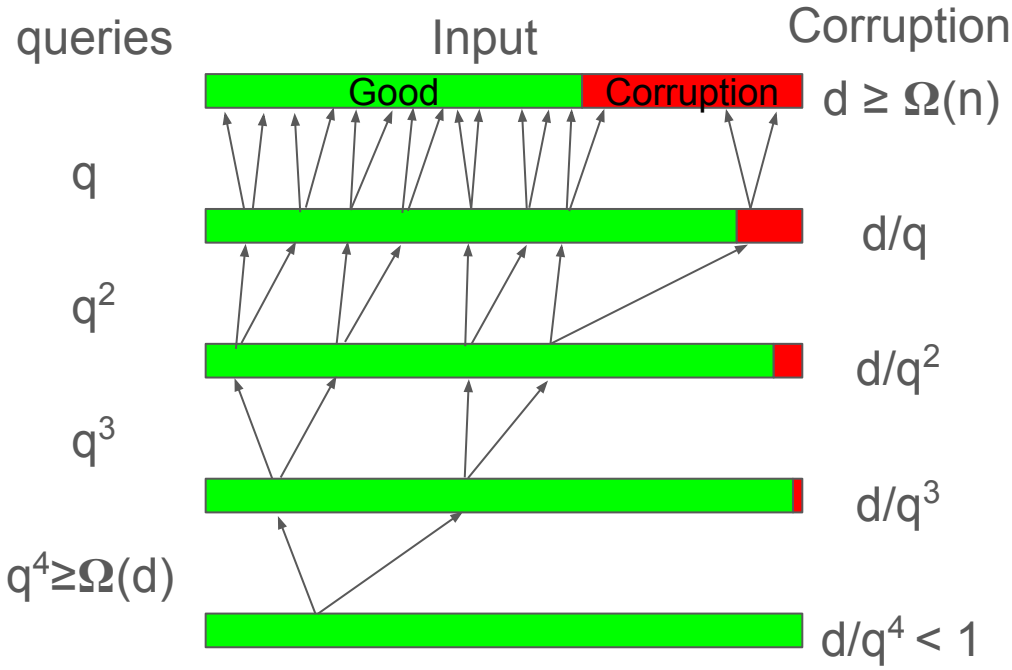

Fix: More than one function

What about m different q query functions $f_1, \ldots, f_m$?

Now the $O(\boldsymbol{\delta} / q)$ failures are distributed among m functions.

Less than $O(\boldsymbol{\delta} / m)$ on average.

Don't have to pay for m in the recursion, so can make $m \gg q$!

Definition (Improving Set)**:**

$f_1, \ldots, f_m$ is a below $\boldsymbol{\delta}$, factor η improving set if for any w and C(x) with $\Pr_i[C(x)_i \neq w_i] \leq \boldsymbol{\delta}$ we have

$$\Pr_{i,j}[C(x)_i \neq f_j(w)_i] \leq \eta \Pr_i[C(x)_i \neq w_i]$$

Example

Try all $f_1, .., f_m$.

Keep going until error is zero.

Still recursive, but fewer levels if $q \ll m$.

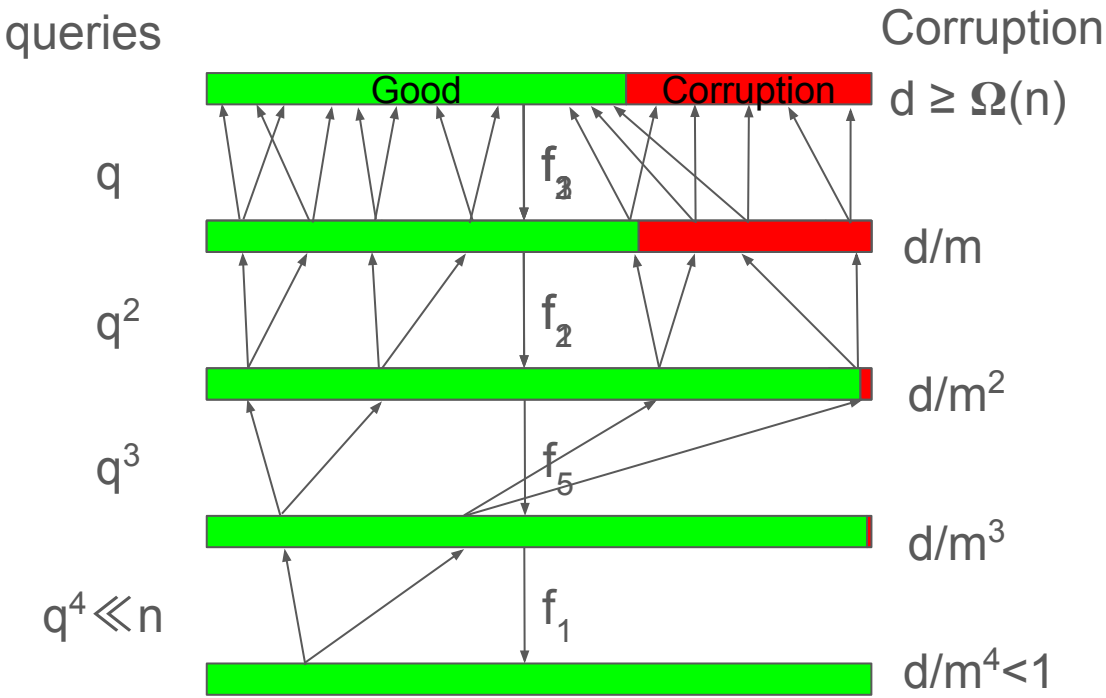

Can We Find such an Improving Set?

Yes, (for typical LCC).

For any q query LCC and η, there is a O(qlog(n)) query improving set with size

$$m = O(\log(n)^2/\eta).$$

If $q = n^{o(1)}$, then setting η appropriately gives space $n^{o(1)}$ time $n^{1+o(1)}$.

Uniform Decoding for Reed-Muller

Second Result: Efficient **Uniform** Decoders

Theorem (Deterministic, Uniform Decoders)**:**

There is a good code with a **uniform**, deterministic decoder running in space $n^{o(1)}$ and time $n^{1+o(1)}$.

The code is based on Lifted Reed-Solomon codes.

Also applies to the specific case of Reed-Muller codes.

Samplers

Family of subsets of N, $\mathscr{S}$.

We say $\mathscr{S}$ is a sampler if:

For all sets A (let $\mu = {}^{|A|}/_{|N|}$).

The probability $S \in \mathscr{S}$ oversamples A is low.

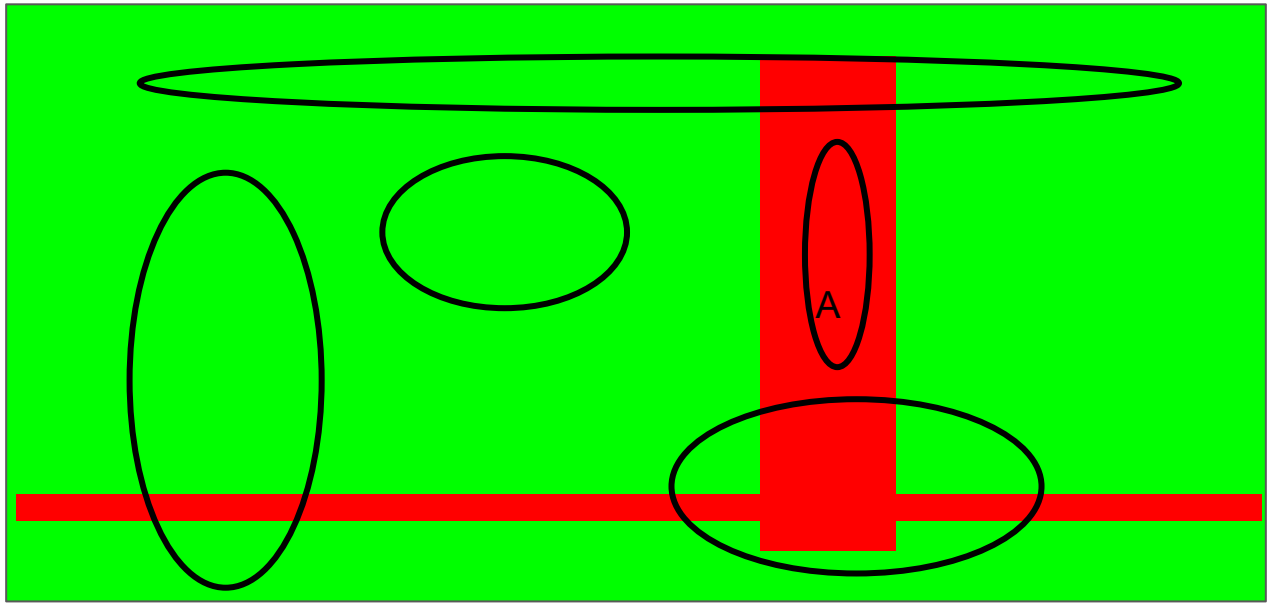

Definition (Sampler): $\mathscr{S} = (S_1,...S_k)$ where $S_i \subseteq N$ is a sampler if for some accuracy error $\varepsilon > 0$ and strong confidence error δ, for all $A \subseteq N$, and $\mu = {}^{|A|}/_{|N|}$ we have

$$\Pr_i[{}^{|S_i \cap A|}/_{|S_i|} \geq \mu + \varepsilon] \leq \delta\mu.$$

Curve Samplers

Need samplers with special structure to allow decoding.

- Lines (Line samplers).
- Subspaces (Space Samplers).
- Curves (Curve Samplers).

Prior curve samplers by Ta-Shma and Umans (and later by Guo) exist, but they:

- Had too many samples, more than n^4, while we need $n^{1+o(1)}$.
- Only proved a weaker notion of confidence error.

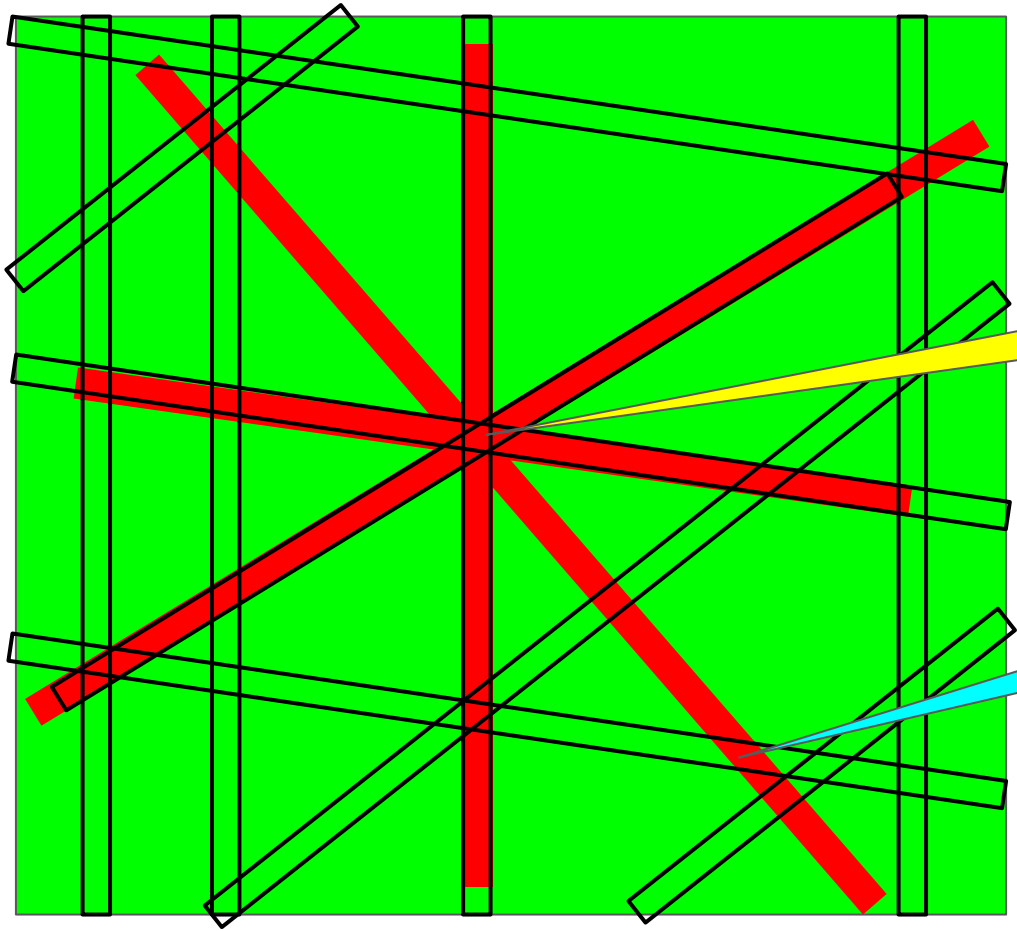

Line Samplers

A sampler whose samples are lines.

Some points still won't be corrected, but few of these

Most points (even corrupted points), most lines through that point don't oversample A.

How good are line samplers?

For q queries, the probability they oversample is about $\eta \cong 1/q$.

Comes from pairwise independence.

NOT GOOD ENOUGH! Need $\eta << 1/q$.

This is the best lines (or subspaces) can do!

Solutions?

- Use curves (works, but gives much worse rate).
- Use several lines through a point
- (extends to lifted Reed-Solomon codes).

Third Result: New Curve Samplers

Theorem (Curve Sampler)**:**

For appropriate degree t, dimension d, and any $\varepsilon > 0$ there is a degree t-curve sampler for $\mathbb{F}^d$ of size n $|\mathbb{F}|^{poly(t)}$ with accuracy error ε and strong confidence error:

$$\delta = O_t\left(1/(\varepsilon|\mathbb{F}|)^{\square(t)}\right)$$

The number of samples is close to $n = |\mathbb{F}|^d$ when $t \ll d$.

The number of queries is $q = |\mathbb{F}|$.

For large t, confidence error is less than 1/q.

Sampler Construction

First sample a subspace.

Epsilon biased sets in extension field. Gives a line sampler, which is a subspace sampler over original field.

Sub-sample with curves.

Since subspace is small, use all low degree curves as a sampler.

Choose **one** curve.

Do low degree correction on that curve.

Or a sample a few lines through a point in that subspace.

Choose **a few** lines.

Correct on each line and take a majority.

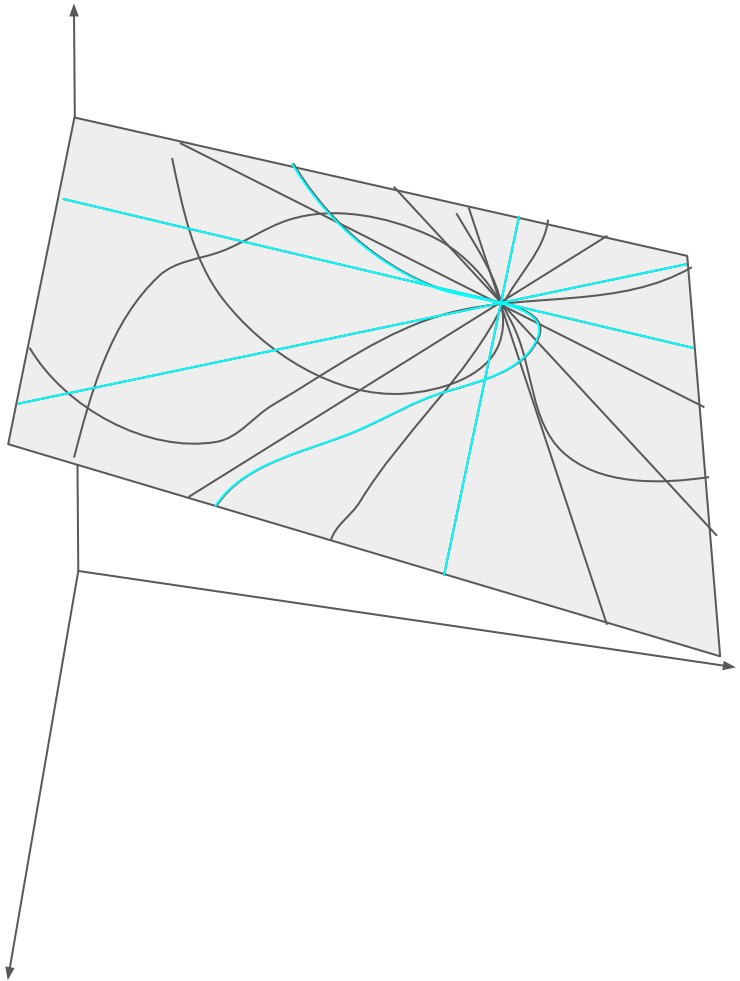

Some Technical Notes

Subspace Samplers

Subspace samplers come from line samplers in extension fields.

- Line samplers through ε-biased sets.
- ε-biased sets (Jalan, Moshkovitz using techniques from Ta-Shma).
- Lines through extension field are subspaces of base field.

Prior curve samplers also use extension fields (Ta-Shma and Umans).

- Uses curves, not lines over extension field.
  - Doesn't let us subsample lines (needed for good rate).
  - Not as randomness efficient ($\Omega(n^2)$ samples).

Getting Better Rate with Fewer Queries.

Reed-Muller gives bad rate for few queries.

Use a closely related code called Lifted Reed-Solomon.

- Low degree only when restricted to lines.
  - Can’t use curve samplers, we use samplers with many lines.
- For high degree and low characteristic, more general than Reed-Muller.

Lifted Reed-Solomon with high rate (and few queries) has low distance.

Use distance amplification: Kopparty, Saraf, and Yekhanin.

Open Problems

1. Find a single code that is encodable and decodable in space $n^{o(1)}$, time $n^{1+o(1)}$.

2. ~~Extend to list decoding.~~
    a. Already did this in a follow up work *for Reed-Muller* codes.
        i. Constants are huge, 10^{20}. Give better constants.
    b. Doesn't have constant rate when achieving space $n^{o(1)}$, time $n^{1+o(1)}$.

3. Give time/space efficient **uniform** decoders for more codes (like multiplicity codes).
    a. Only gave non-uniform decoders for multiplicity codes.

4. Find a deterministic decoder running in space polylog(n) and time n polylog(n).

Thanks for Listening