CS250/EE387 - LECTURE 7 - Efficiently decoding concatenated codes.

AGENDA

1. ZYABLOV BOUND
2. EFFICIENTLY' DECODING CONCATENATED CODES

Recall the GOAL from last lecture:

GOAL. Obtain EXPLICIT (aka, efficiently constructible), ASYMPTOTICALLY GOOD families of BINARY CODES, ideally with fast algorithms.

(for today

Today, we'll see how to use CONCATENATED CODES to achieve this goal.

RECALL:

The CONCATENATED CODE $C_{in} \circ C_{out} \subseteq \sum_{in}^{n_{out} \cdot n_{in}}$ is defined [by picture...] by

$X \in \sum_{in}^{k_{in} \cdot k_{out}} \sim \sum_{out}^{k_{out}}$

$\alpha \in \sum_{out}$

$c \in \sum_{out}^{n_{out}}$, encoding of x under C_{out}.

$c' \in \sum_{in}^{n_{in}}$, encoding of α under C_{in}.
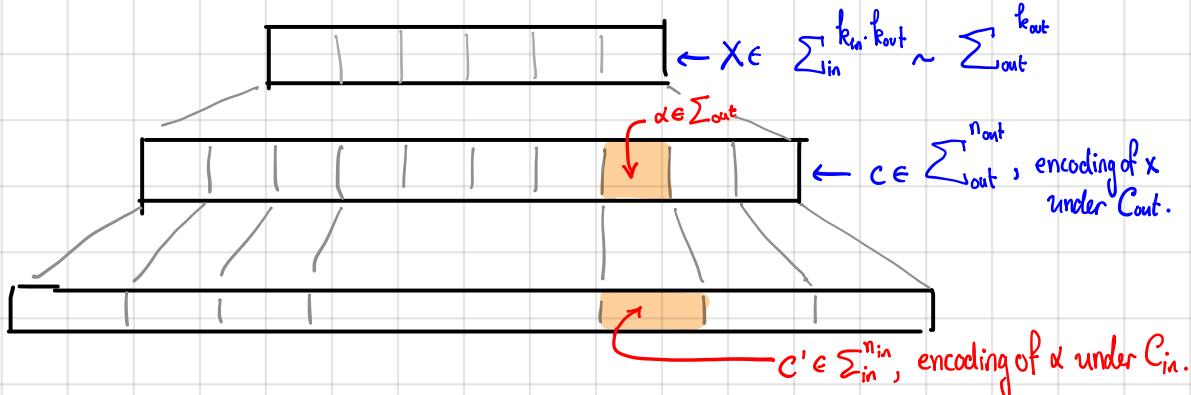
① ZYABLOV BOUND

THM. For any $\varepsilon > 0$, there is a family of explicit binary linear codes of rate R and distance δ, (at least) for any R,δ satisfying

$$R \geq \sup_{0 < r < 1 - H_2(\delta) - \varepsilon} \left(r \cdot \left(1 - \frac{\delta}{H_2^{-1}(1 - r - \varepsilon)} \right) \right)$$

That is called the <u>Zyablov Bound</u>.

Best EXPLICIT binary code lives somewhere in here.
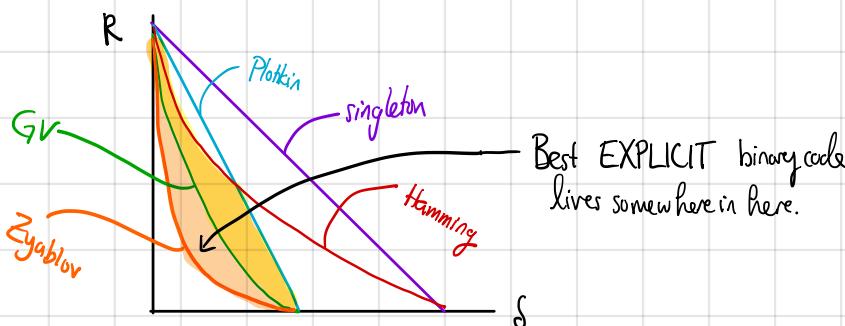
pf. The construction will be:

- C_{out} = RS code with rate R_{out}, dist. $\delta_{out} \geq 1 - R_{out}$
- C_{in} = <u>Binary linear code on the GV bound</u>, with rate $r \geq 1 - H_2(\delta_{in}) - \varepsilon$.

We still need to say how to get this... we'll get there.

- The concatenated code has:

RATE: $R_{out} \cdot r$

DISTANCE: $\delta_{out} \cdot \delta_{in}$ (at least) $= (1 - R_{out}) H_2^{-1}(1 - r - \varepsilon)$.

Hence $R = R_{out} \cdot r = r \cdot \left(1 - \frac{\delta_{in} \cdot \delta_{out}}{H_2^{-1}(1 - r - \varepsilon)} \right)$,

(solve for R_{out} in the above)

and then we get to choose r.

How to choose these parameters? Given δ and $r \leq 1 - H_2(\delta) - \varepsilon$, we want $\delta_{in} \cdot \delta_{out} \geq \delta$, and $r \geq 1 - H_2(\delta_{in}) - \varepsilon$. So choose $\delta_{in} \in (0, \frac{1}{2})$ so that $\delta_{in} = H_2^{-1}(1 - r - \varepsilon)$ and then choose $\delta_{out} = \delta / \delta_{in}$. For that to make sense we need $\delta_{in} > \delta$, but fortunately we restricted $r < 1 - H_2(\delta) - \varepsilon$ in the thm statement, aka $\delta < H_2^{-1}(1 - r - \varepsilon) = \delta_{in}$.

So that gives us the combinatorial bound.
Next, the algorithmic bit.

proof continued ...

Suppose the evaluation pts for the RS code are $\mathbb{F}_q^*$, where $q = 2^{k_{in}}$.

So $k_{in} = \lg(q)$, and $n_{out} = q - 1$

ALG 1. Search over all $\mathbb{F}_2$-linear codes of rate r and dimension k_{in}.

There are approximately $2^{n_{in} \cdot k_{in}} = 2^{k_{in}^2/r} = 2^{\lg^2(q)/r}$ such codes

$$q = n_{out} + 1 = \frac{n}{n_{in}} + 1 = \frac{r\,n}{k_{in}} + 1 \Rightarrow n \sim \frac{1}{r} q \lg(q)$$

↑ final block length

So $\lg^2(q) = \Theta(\lg^2(n))$, so that's $2^{\Theta(\lg^2(n))} = n^{\Theta(\lg(n))}$, which is **NOT** polynomial time. :(

ALG 2.   In class, we will give an alg to construct binary linear codes on the GV bound w/ rate r, dim k_{in} in time $2^{O_r(k_{in})}$, instead of $2^{O_r(k_{in}^2)}$! This will fix the above, and proves the theorem.

So we have acheived part of our goal.   (Explicit codes — no algs yet).

HOWEVER, this version of "explicit" [can compute it in polynomial time] may be unsatisfying.

WHAT IF I WANTED "explicit" meaning: "Give me a short, useful description"
    Formally, I'd like to be able to compute any entry G_{ij} in time polylog(n).

__IDEA__: Instead of using the same inner code at every position and requiring it to be good, we'll use a different inner code in each position.

We won't actually know which of these inner codes is good, but we'll know that enough of them are good.

__THM.__ Let $\varepsilon > 0$, fix any k. There is an ensemble of binary linear codes

$$C_{in}^1, C_{in}^2, \dots, C_{in}^N \subseteq \mathbb{F}_2^{2k}$$

of rate $\frac{1}{2}$, with $N = 2^k - 1$, so that for at least $(1-\varepsilon)N$ values of i, C_{in}^i has distance at least $H_2^{-1}(\frac{1}{2} - \varepsilon)$.

This is called the __WOZENCRAFT ENSEMBLE__.

proof idea. For $x \in \mathbb{F}_2^k$, treat it as an element of $\mathbb{F}_{2^k}$. Then for each $\alpha \in \mathbb{F}_{2^k}^*$, let the α^{th} code C_{in}^α be the image of the encoding map

$$E_{in}^\alpha : x \longmapsto (x, \alpha \cdot x)$$

multiplication in $\mathbb{F}_{2^k}$

treat these as $2k$ bits.

FUN EXERCISE: finish the proof!

Using the Wozencraft ensemble, we can implement the idea above to obtain the __JUSTESAN CODE__.

DEF (JUSTESEN CODE)

Let $k > 0$ [k will be the dimension of the inner codes in the Wozencraft Ensemble]

Let $C_{out} = RS_{2^k}\left(\mathbb{F}_{2^k}^*, \ 2^k-1, \ R_{out}\cdot(2^k-1)\right)$

Use the Wozencraft Ensemble as the inner code:

$$C = \left\{ \ \left\langle E_{in}^{\alpha}\left(f(\alpha)\right) \right\rangle_{\alpha \in \mathbb{F}_{2^k}^*} \ : \ f \in \mathbb{F}_{2^k}[X], \ deg(f) < R_{out}(2^k-1) \right\}$$

CLAIM. Let R_{out} be constant. Then C is asymptotically good!

pf.
(sketch) The rate is $R_{out}/2$, and it's a binary linear code, so we just have to consider the minimum wt to compute the distance. Choose $\varepsilon > \frac{1-R_{out}}{2}$. Consider any codeword:
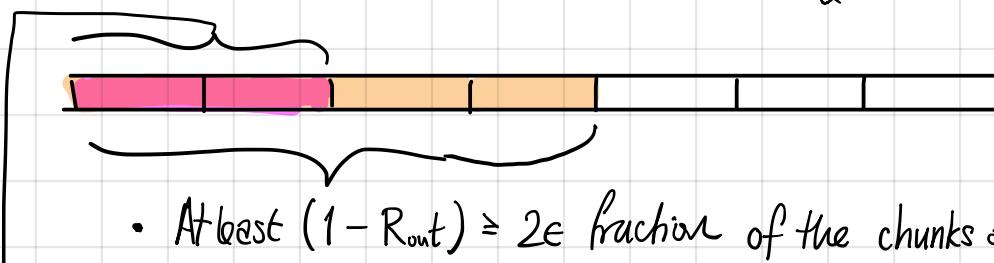
- At least $(1-R_{out}) \geq 2\varepsilon$ fraction of the chunks are the encodings of nonzero symbols.

- At most an ε-fraction of chunks have "bad" inner codes, so at least an $2\varepsilon - \varepsilon = \varepsilon$-fraction of chunks are the encodings of nonzero symbols with a "good" inner code.

For each of those, since the inner code has distance $\geq H_2^{-1}(\frac{1}{2}-\varepsilon) = \Theta(1)$, a constant fraction of the bits in each of a constant fraction of blocks are nonzero.

$\Rightarrow$ Each nonzero codeword has relative weight larger than some constant. Thus the code is asymptotically good.

So the JUSTESEN CODE is "EXPLICIT" in the way we wanted.
The α'^{th} block is given by $(f(\alpha), \alpha \cdot f(\alpha)) \in \mathbb{F}_q^2 \rightsquigarrow \mathbb{F}_2^{2\lg(q)}$.
That's pretty explicit!

FUN EXERCISE. What is the best rate/distance trade-off you can get w/ the
Justesen code?

FUN EXERCISE. What happens to the Wozencraft ensemble if you do
$$x \longmapsto (x, \alpha \cdot x, \alpha^2 \cdot x, ..., \alpha^r \cdot x) \ ?$$

\END{Aside}

This space intentionally blank
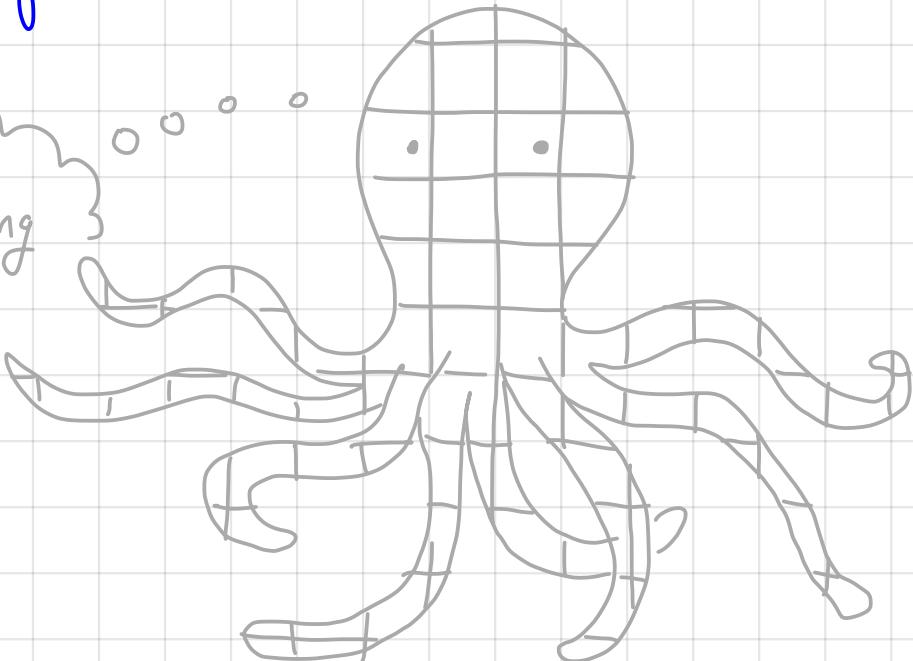
② EFFICIENT DECODING ALGS for CONCATENATED CODES.

Now that we have explicit constructions of asymptotically good
codes (and in particular efficient encoding algs),
what about efficient DECODING algs?

FIRST TRY at decoding:

$x \in \sum_{out}^{k_{out}} \sim \sum_{in}^{k_{in} \cdot k_{out}}$

C_{out}

$\alpha \in \sum_{out}$

$C_{in} \quad C_{in} \quad C_{in} \qquad C_{in} \quad C_{in} \quad C_{in}$

$c \in C_{in} \circ C_{out} \longrightarrow$

$c' \in \sum_{in}^{n_{in}}$, encoding of α under C_{in}.

ADD ERROR

$\tilde{c} \in \sum_{in}^{n_{in} \cdot n_{out}}$
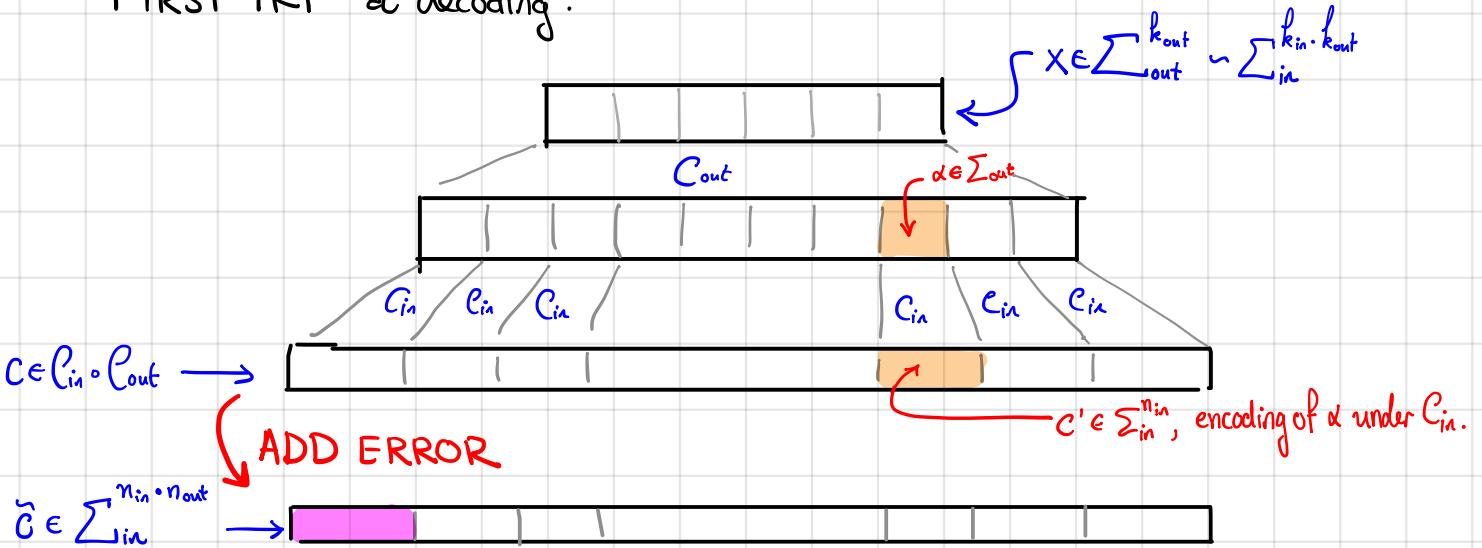
① Decode each of these ▮ blocks: that is, find the codeword $c' \in C_{in}$ which is the closest to the received word.

② Convert the "corrected" chunks ▭ $\in C_{in}$ into ▣ $\alpha \in \sum_{out}$

③ Decode C_{out} to get the original message.

CLAIM.*   The above works PROVIDED that the number of errors e is $< \dfrac{d_{in} \cdot d_{out}}{4}$

* Maybe with some floors and/or ±1's.

NOTICE: $d = d_{in} \cdot d_{out}$ is the designed distance of the concatenated code.
So we'd really like $e \le \left\lfloor \dfrac{d-1}{2} \right\rfloor$, not $d/4$.
But let's prove the claim anyway, to understand why this approach might fail.

pf (ish).   Let's call a block ▮ "BAD" if there are more than $\left\lfloor \dfrac{d_{in}-1}{2} \right\rfloor$ errors in that block.
If there are e errors total, at most $e / \left\lfloor \dfrac{d_{in}-1}{2} \right\rfloor$ blocks are BAD.
If a block is NOT BAD, then the inner code works.
Thus we win provided (#BAD BLOCKS) $\le \left\lfloor \dfrac{d_{out}-1}{2} \right\rfloor$

aka $\quad e / \left\lfloor \dfrac{d_{in}-1}{2} \right\rfloor \le \left\lfloor \dfrac{d_{out}-1}{2} \right\rfloor$

$e \le \left\lfloor \dfrac{d_{in}-1}{2} \right\rfloor \left\lfloor \dfrac{d_{out}-1}{2} \right\rfloor \approx \dfrac{d_{in} \cdot d_{out}}{4}$

Indeed, that's what happens when there are exactly $\left\lfloor \dfrac{d_{in}-1}{2} \right\rfloor$ errors in each bad block.

The proof shows that this might NOT be a good idea.

If the adversary JUST BARELY messes up as many blocks as they can, this decoder will fail on $\lfloor \frac{d-1}{2} \rfloor$ errors.

WHAT ARE WE LEAVING ·ON THE TABLE?

Key observation: When we decode the inner code $\boxed{}_{\hat{c} \in \Sigma_{in}^{n_{in}}} \rightarrow \boxed{}_{c' \in C_{in}}$,

we learn more than just $\boxed{}_{c' \in C_{in}}$; we also know $wt\left(\boxed{}_{\hat{c} \in \Sigma_{in}^{n_{in}}} - \boxed{}_{c' \in C_{in}} \right)$.

SOME MOTIVATING EXAMPLES:

① Each block either has 0 or $d_{in}/2$ errors. [This is the bad example from before].

$\boxed{\frac{d_{in}}{2} \text{ errors in each of } \binom{e}{\frac{d_{in}}{2}} \text{ blocks}}$

C

$\downarrow$

$c \in C$

$\downarrow$ Correct each block w/ C_{in}

$\hat{c}$

← These blocks have no errors and don't change when we decode them.

↑ This block had some errors. When we decode it, it's to something at least $d_{in}/2$ away, because:

$\geq d_{in}$

$\leq \frac{d_{in}}{2}$

this distance is also $\geq d_{in}/2$.
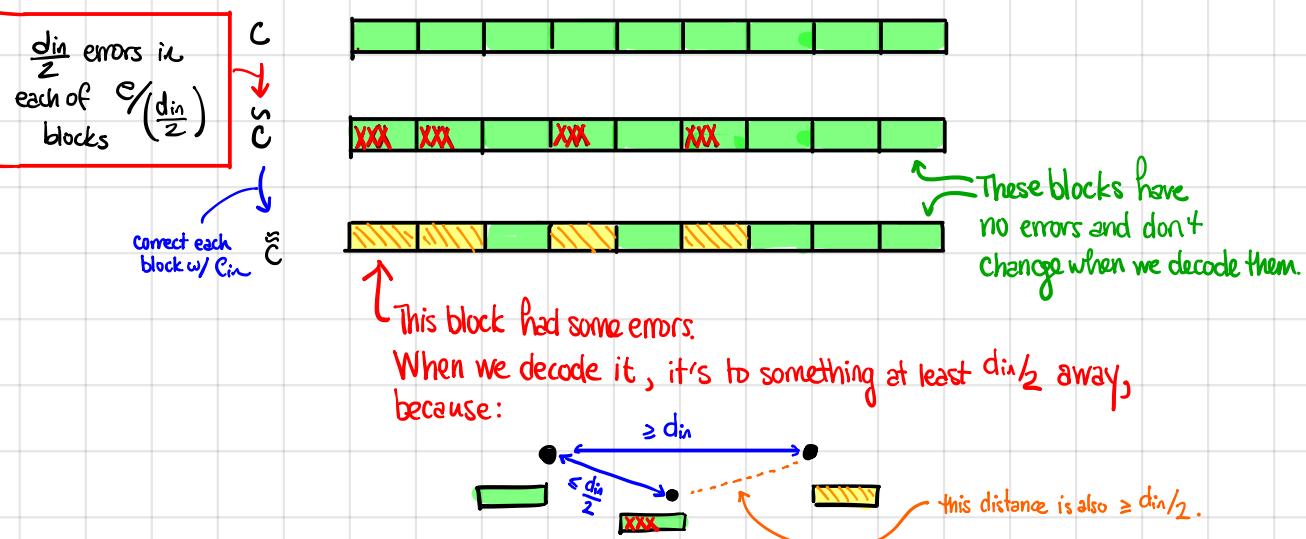
Thus, even though the $\boxed{}$ blocks are incorrect, we can detect that they were incorrect.

So the thing we should do in this case is treat the $\boxed{}$ blocks as ERASURES. We can handle twice as many of those! So our error tolerance is actually about $d/2$ in this case, which is what we wanted.

② MOTIVATING EXAMPLE #2. The bad guy tries to foil our previous example by adding error d_{in} to some blocks, turning them into other codewords.

d_{in} errors in each of e/d_{in} blocks

correct each block. They are all already codewords! Nothing left to do.
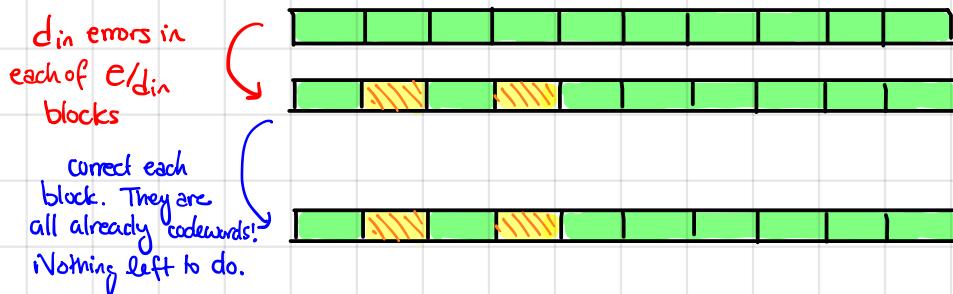
Now we can't detect anything! BUT, there are only e/d_{in} corrupted blocks. Again we save a factor of 2 and can correct up to $e \approx d/2$ errors.

We would like to interpolate between these two extremes.

CLAIM. We can efficiently decode $RS_q(n,k)$ from e errors and s erasures, as long as $2e + s < n - k + 1$. (aka, the distance of the RS code).

pf-ish. Throw out the s erasures. You are left with $RS(n-s, k)$.
Since $2e < (n-s) - k + 1$, run Berlekamp-Welch to correct the errors. Now you have an $RS(n,k)$ codeword w/ s erasures. Since $s < n-k+1$, correct the erasures (via linear algebra).

This inspires the following algorithm:

ALGORITHM: (NOT THE FINAL VERSION).

Given $\vec{w} = (w_1, w_2, \ldots, w_{n_{out}}) \in \left(\mathbb{F}_{q_{in}}^{n_{in}} \right)^{n_{out}}$, s.t. $\Delta(w, c) < \dfrac{d_{in} \cdot d_{out}}{2}$ for some $c \in C_{in} \circ C_{out}$

For each $i = 1, \ldots, n_{out}$:

$\quad$ Let $w_i' = \underset{y \in C_{in}}{\text{argmin}} (\Delta(y, w_i))$

$\quad$ With probability $\min\left(\dfrac{2\Delta(w_i, w_i')}{d_{in}}, 1 \right)$:

$\quad\quad$ set $\beta_i = \perp$

$\quad$ Else:

$\quad\quad$ Set β_i s.t. $E_{in}(\beta_i) = w_i'$

Run C_{out}'s (error + erasure) decoder on $(\beta_1, \ldots, \beta_{n_{out}})$, RETURN the result.

Why does this algorithm work?

CLAIM. $\mathbb{E}_{\substack{\text{algorithm's} \\ \text{randomness}}}\left[\left(\#\beta_i \text{ that} = \perp\right) + 2\cdot\left(\#\beta_i \substack{\text{ that are not} \\ \text{correct}}\right)\right] < d_{out}.$

Proof (sketch). Let $e_i = \Delta(w_i, c_i)$, so
$$e = \sum_i e_i < \frac{d_{in}\cdot d_{out}}{2}.$$

Let $X_i^{\perp} = \mathbb{1}\{\beta_i = \perp\}$
$X_i^{e} = \mathbb{1}\{\beta_i \neq \perp \text{ and } \beta_i \neq \alpha_i\}$

SUB CLAIM. $\mathbb{E}\left[2X_i^{e} + X_i^{\perp}\right] \leq \frac{2e_i}{d_{in}}.$

(Notice that the SUBCLAIM proves the CLAIM, by linearity of expectation).

pf. of SUBCLAIM:

CASE 1. $C_i = w_i'$. So $X_i^e = 0$, and $\mathbb{E}\left[X_i^{\perp}\right] \leq \frac{2\cdot\Delta(w_i, w_i')}{d_{in}} = \frac{2\Delta(w_i,c_i)}{d_{in}} = \frac{2e_i}{d_{in}}$

CASE 2. $C_i \neq w_i'$ As before, $\mathbb{E}\left[X_i^{\perp}\right] = \frac{\min(2\cdot\Delta(w_i,w_i'), d_{in})}{d_{in}}$, and

$\mathbb{E}\left[X_i^{e}\right] = 1 - \mathbb{E}\left[X_i^{\perp}\right]$, since if we didn't find $\perp$, then we made a mistake.

SUBSUBCLAIM. $2e_i + \min\left(2\Delta(w_i, w_i'), d_{in}\right) \geq 2d_{in}.$

Given the $\substack{\text{SUB} \\ \text{SUB}}$CLAIM, $\mathbb{E}\left[2X_i^e + X_i^{\perp}\right] = 2(1 - \mathbb{E}[X_i^{\perp}]) + \mathbb{E}[X_i^{\perp}]$
$$= 2 - \mathbb{E}[X_i^{\perp}]$$
$$\xrightarrow{\text{SUBSUBCLAIM}} \leq 2 - \frac{1}{d_{in}}\min\left(2\cdot\Delta(w,w_i'), d_{in}\right)$$

(Proof of SUBSUBCLAIM on next page).
$$\leq 2 - \frac{1}{d_{in}}\left(2d_{in} - 2e_i\right)$$
$$= 2e_i/d_{in}, \text{ as desired.}$$
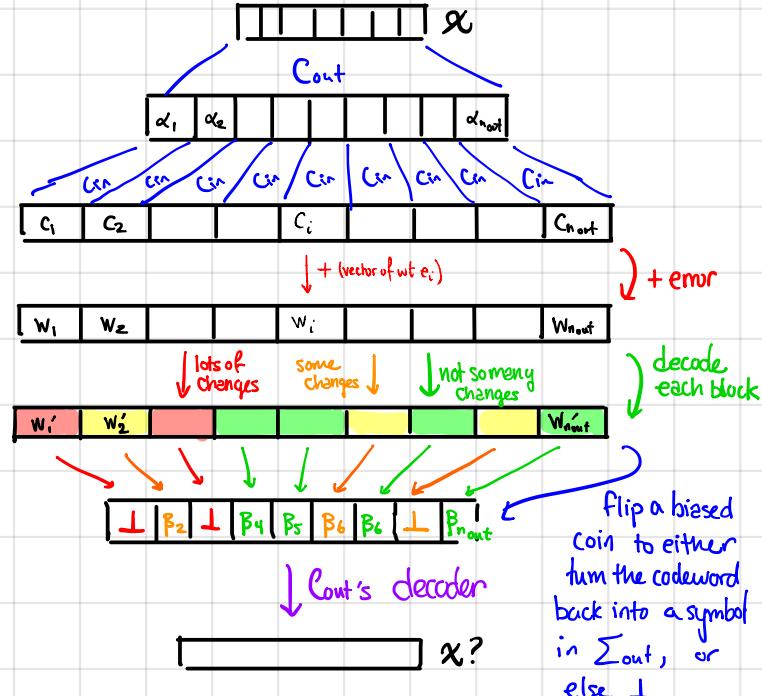
If $c_i \neq w_i'$, $\quad 2e_i + \min\left(2\Delta(w_i, w_i'), d_{in}\right) \geq 2 \cdot d_{in}$

Proof: Suppose that $2\Delta(w_i, w_i') \leq d_{in}$. Then the SUBSUBCLAIM reads:

$$2e_i + 2\Delta(w_i, w_i') \geq 2d_{in}$$

$$e_i + \Delta(w_i, w_i') \geq d_{in}$$

$$\Delta(w_i, c_i) + \Delta(w_i, w_i') \geq d_{in}$$

which is true since

$$d_{in} \leq \Delta(c_i, w_i') \leq \Delta(w_i, c_i) + \Delta(w_i, w_i')$$

↑ Since $c_i, w_i' \in C_{in}$ and $c_i \neq w_i$

↑ triangle inequality.

On the other hand, if $d_{in} < 2\Delta(w_i, w_i')$, then the SUBSUBCLAIM reads:

$$2e_i + d_{in} \geq 2d_{in} \quad \text{aka} \quad e_i \geq \frac{d_{in}}{2}.$$

But this must be true because we are in the setting where $c_i \neq w_i'$. Indeed, if $e_i < \frac{d_{in}}{2}$, then the inner code's decoder would have worked correctly and we would have $c_i = w_i'$.

So the CLAIM implies that the algorithm works "in expectation."

We could try to turn this into a high probability result (repeat a bunch of times), but instead we will actually be able to DERANDOMIZE it.

STEP 1.   We will reduce the necessary randomness by a little bit.

ALGORITHM VERSION 2

Given $\vec{w} = (w_1, w_2, \ldots, w_{n_{out}}) \in \left(\mathbb{F}_{q_{in}}^{n_{in}}\right)^{n_{out}}$, s.t. $\Delta(w, c) < \dfrac{d_{in} \cdot d_{out}}{2}$
for some $c \in C_{in} \circ C_{out}$

CHOOSE $\theta \in [0, 1]$ UNIFORMLY AT RANDOM.

For each $i = 1, \ldots, n_{out}$:

$\quad$ Let $w_i' = \underset{y \in C_{in}}{\text{argmin}} \, (\Delta(y, w_i))$

$\quad$ IF $\theta \le \min\left(\dfrac{2\Delta(w_i, w_i')}{d_{in}}, 1\right)$:

$\quad\quad$ set $\beta_i = \perp$

$\quad$ Else:

$\quad\quad$ Set β_i s.t. $E_{in}(y_c) = w_i'$

Run C_{out}'s (error + erasure) decoder on $(\beta_1, \ldots, \beta_{n_{out}})$, RETURN the result.

That is, we never used the fact that our draws for β_i were independent. So let's make them not at all independent.

Our next step will be to search over all possible θ's. In fact, we only need to look at $n_{out} + 2$ values of θ:
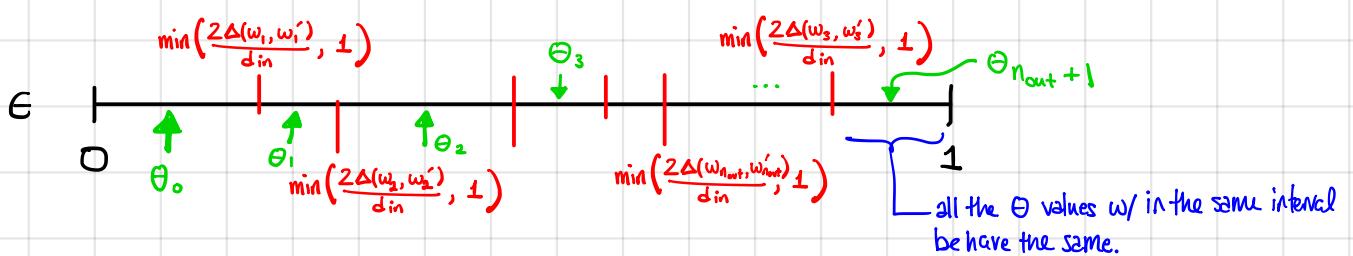
ALGORITHM: FINAL VERSION

Given $\vec{w} = (w_1, w_2, \ldots, w_{n_{out}}) \in \left(\mathbb{F}_{q_{in}}^{n_{in}}\right)^{n_{out}}$, s.t. $\Delta(w, c) < \dfrac{d_{in} \cdot d_{out}}{2}$

for some $c \in C_{in} \circ C_{out}$

COMPUTE THE $n_{out} + 2$ RELEVANT VALUES of $\Theta, \Theta_0, \ldots, \Theta_{n_{out}+1}$

FOR $j = 0, \ldots, n_{out} + 1$:

For each $i = 1, \ldots, n_{out}$:

Let $w_i' = \underset{y \in C_{in}}{\mathrm{argmin}}\left(\Delta(y, w_i)\right)$

IF $\Theta_j < \min\left(\dfrac{2\Delta(w_i, w_i')}{d_{in}}, 1\right)$:

set $\beta_i = \perp$

Else:

Set β_i s.t. $E_{in}(y_i) = w_i'$

Run C_{out}'s (error + erasure) decoder on $(\beta_1, \ldots, \beta_{n_{out}})$, **to obtain $\breve{x}$**

IF $\Delta(E_{NC}(\breve{x}), w) \leq \left\lfloor \dfrac{d-1}{2} \right\rfloor$:

RETURN $\breve{x}$

The fact that this algorithm is correct follows from our earlier claim.

Since $\underset{\Theta}{\mathbb{E}}\left[2 \cdot (\#errs) + (\#erasures)\right] \leq d_{out}$,

there exists some $\Theta \in [0, 1]$ so that $2(\#errs) + (\#erasures) \leq d_{out}$, aka so that the alg. finds the correct $\breve{x}$.

Thus, our algorithm above, which tries ALL values of Θ, must find that good value and return the correct answer.

What is the running time of this algorithm?

Depends on the codes. Let's choose our explicit construction

$\downarrow$

Recall we had $n_{out} = q_{out} - 1$, and $q_{out} = 2^{k_{in}}$.

- C_{out} = RS code with rate R_{out}, dist. $\delta_{out} = 1 - R_{out}$
- C_{in} = Binary linear code on the GV bound, with rate $r \geq 1 - H_2(\delta_{in}) - \epsilon$.
  You showed/will show how to find this in time poly(n) on your homework.

The expensive bits of the alg are:

For $O(n_{out})$ choices of Θ:
  for $i = 1, ..., n_{out}$:
   - Decode the inner code (length $n_{in} = O(k_{in}) = O(\log(n_{out}))$) by brute force // Time $O(n_{in} \cdot |C_{in}|) = O(n_{in} \cdot 2^{k_{in}}) = poly(n)$

   - Run the RS decoder.   // Time poly(n)

So altogether the whole thing runs in polynomial time.

We have proved

THM   For every $R \in (0,1)$, there is a family C of EXPLICIT BINARY LINEAR CODES that lies at or above the Zyablov bound. Further, C can be decoded from errors up to half the Zyablov bound in time poly(n).

AKA, we have achieved our goal! Hooray!

To RECAP the story of Concatenated Codes:

- We considered (RS code) ∘ (Binary Linear Code on the GV bound)

- Because the inner code is so small, we can find a good one by brute force in time poly(n).

- We can be a little more clever with the Justesen Code, if we want something asymptotically good and STRONGLY explicit.

- (RS) ∘ (Binary code on the GV bd) met the "Zyablov Bound", which was defined as "the bound that these codes meet."

- We saw how to use Forney's GMD decoder to efficiently decode these codes up to half the minimum distance.

QUESTIONS to PONDER :

① When does code concatenation give distance STRICTLY LARGER than $d_{in} \cdot d_{out}$?

② Do there exist concatenated codes on the GV bound?
   SPOILER ALERT: YES, see [Thomesson 1983].   (It's a randomized construction)

③ Can we decode these ↗ efficiently?
   SPOILER ALERT: ALSO YES.   It uses list decoding, we may see it later.

④ Can you do better than the Zyablov bound for EXPLICIT CODES with EFFICIENT ALGS?